

The PHOBOS Perspective on Discoveries at RHIC

B.B.Back,¹ M.D.Baker,² M.Ballintijn,⁴ D.S.Barton,² B.Becker,² R.R.Betts,⁶ A.A.Bickley,⁷ R.Bindel,⁷
A.Budzanowski,³ W.Busza,⁴ A.Carroll,² Z.Chai,² M.P.Decowski,⁴ E.García,⁶ T.Gburek,³ N.K.George,^{1,2}
K.Gulbrandsen,⁴ S.Gushue,² C.Halliwell,⁶ J.Hamblen,⁸ A.S.Harrington,⁸ M.Hauer,² G.A.Heintzelman,²
C.Henderson,⁴ D.J.Hofman,⁶ R.S.Hollis,⁶ R.Hołyński,³ B.Holzman,^{2,6} A.Iordanova,⁶ E.Johnson,⁸ J.L.Kane,⁴
J.Katzy,^{4,6} N.Khan,⁸ W.Kucewicz,⁶ P.Kulinich,⁴ C.M.Kuo,⁵ J.W.Lee,⁴ W.T.Lin,⁵ S.Manly,⁸ D.McLeod,⁶
A.C.Mignerey,⁷ M.Nguyen,² R.Nouicer,^{2,6} A.Olszewski,³ R.Pak,² I.C.Park,⁸ H.Pernegger,⁴ C.Reed,⁴
L.P.Remsberg,² M.Reuter,⁶ C.Roland,⁴ G.Roland,⁴ L.Rosenberg,⁴ J.Sagerer,⁶ P.Sarin,⁴ P.Sawicki,³ H.Seals,²
I.Sedykh,² W.Skulski,⁸ C.E.Smith,⁶ M.A.Stankiewicz,² P.Steinberg,² G.S.F.Stephans,⁴ A.Sukhanov,²
J.-L.Tang,⁵ M.B.Tonjes,⁷ A.Trzupek,³ C.M.Vale,⁴ G.J.van Nieuwenhuizen,⁴ S.S.Vaurnovich,⁴ R.Verdier,⁴
G.I.Verés,⁴ E.Wenger,⁴ F.L.H.Wolfs,⁸ B.Wosiek,³ K.Woźniak,³ A.H.Wuosmaa,¹ B.Wysłouch,⁴ and J.Zhang⁴
(PHOBOS collaboration)

¹Argonne National Laboratory, Argonne, IL 60439-4843, USA

²Brookhaven National Laboratory, Upton, NY 11973-5000, USA

³Institute of Nuclear Physics PAN, Kraków, Poland

⁴Massachusetts Institute of Technology, Cambridge, MA 02139-4307, USA

⁵National Central University, Chung-Li, Taiwan

⁶University of Illinois at Chicago, Chicago, IL 60607-7059, USA

⁷University of Maryland, College Park, MD 20742, USA

⁸University of Rochester, Rochester, NY 14627, USA

(Dated: September 1, 2004)

Starting in the summer of 2000, the Relativistic Heavy Ion Collider (RHIC) at the Brookhaven National Laboratory has been providing collisions ranging from p+p up to Au+Au at nucleon-nucleon center-of-mass energies up to 200 GeV. The PHOBOS detector has been operational since the beginning of the RHIC experimental program. This paper describes the conclusions that can be drawn from the data taken thus far with the PHOBOS detector. In the most central Au+Au collisions at the highest RHIC energy, evidence is found for the formation of a very high energy density system whose description in terms of simple hadronic degrees of freedom is inappropriate. Furthermore, the constituents of this novel system are found to interact strongly. The properties of particle production at RHIC energies are shown to follow a number of simple scaling behaviors, some of which continue trends found at lower energies or in simpler systems. As a function of centrality, the total number of charged particles scales with the number of participating nucleons. The constant of proportionality for d+Au is close to the total multiplicity seen in p+p, but for Au+Au the constant is closer to the total multiplicity in e^+e^- or in p+p at roughly twice the center-of-mass energy. When comparing Au+Au at different centralities, the dependence of the yield on the number of participants at higher p_T (~ 4 GeV/c) is very similar to that at low transverse momentum. The measured values of charged particle pseudorapidity density and elliptic flow were found to be independent of energy over a broad range of pseudorapidities when viewed in the rest frame of one of the colliding nuclei, a property we describe as “extensive longitudinal scaling”. Finally, the centrality and energy dependences of several observables were found to factorize to a surprising degree.

PACS numbers: 25.75.-q

I. INTRODUCTION

Currently, there exists a good understanding of the basic building blocks of matter, and of the fundamental forces or interactions between them. The bulk of matter is made of quarks. The strong forces that bind quarks into neutrons, protons, and subsequently nuclei are mediated by the field quanta, the gluons. The fundamental interactions between these partons (quarks and gluons) are described by the theory of Quantum Chromodynamics (QCD) [1] and are reasonably well understood. However, because of the non-abelian nature and strength of the interactions, finding solutions to the QCD equations

remains notoriously difficult. As a result, the current understanding of the phase structure of strongly interacting matter (what phases exist, what are the properties of the matter in each phase, and what is the nature of the transitions between phases) is only partly based on QCD. Instead, it is driven, to a large extent, by experiment. Among many examples of the significance of the properties of QCD “matter” is the fact that more than 98% of the mass of all normal hadronic matter in the universe arises from the interactions, not from the (current) mass of the valence quarks in the hadrons [2]. This mass is generated predominantly by the lower energy interactions which are most difficult to study quantitatively.

Areas of impact outside nuclear physics include the evolution of the early universe and the detailed structure of dense stellar remnants. Both theory and experiment suggest the existence of a very rich “condensed matter” of QCD.

At very short distances (\ll hadronic sizes) the QCD coupling constant between partons is weak and decreases as the distance between the partons decreases, a phenomenon known as “asymptotic freedom” [3, 4]. An expected consequence of asymptotic freedom is that a system created by heating the vacuum to high temperatures should have the properties of an almost ideal relativistic gas in which color is deconfined (first pointed out by [5] using the term “quark soup”, see also [6–8]). The high temperature of this medium entails an extremely high concentration of partons, whose thermodynamics follows the Stefan-Boltzmann law. Such a system has traditionally been designated the Quark-Gluon Plasma (QGP). To specifically recognize its ideal, weakly interacting nature, we use the term wQGP. It is generally believed that before about 10 μ sec following the big bang the whole universe was in the wQGP state.

At another extreme, it is known that the only stable configuration of strongly interacting matter at low temperatures and densities is the multitude of varieties of color neutral objects, namely the hadrons, as well as conglomerates of hadrons such as atomic nuclei. In addition, the QCD Lagrangian (and the wQGP solution of that Lagrangian) is understood to have a higher symmetry than the observed hadron states. The solutions of QCD at temperatures and densities which correspond to normal matter, i.e. the world of hadrons and nuclei, spontaneously break the so-called “chiral symmetry” (see, for example, [9, 10]). The questions of what forms and phases of QCD matter exist between the two extremes and what symmetries, properties, and interactions characterize these phases, are currently the subject of very active theoretical and experimental research.

On both the experimental and the theoretical fronts, there are very few tools available for the study of QCD matter as a function of density and temperature. The most fruitful approach to date to the theoretical study of high temperature QCD has been the use of numerical simulations based on the techniques of lattice gauge theory. Lattice gauge simulations suggest that at low baryon densities there is a phase difference in QCD matter below and above a critical temperature $T_c \sim 150$ – 200 MeV or energy density ~ 1 GeV/fm³ (see, for example, [11], which quotes a T_c of 175 MeV and an energy density of 700 MeV/fm³ $\pm 50\%$). At another extreme, theoretical progress has been made in recent years in the understanding of cold, ultra-dense, QCD matter [12–14]. For example, there are indications that a dense, cold system of equal numbers of u, d and s quarks can form a color superconductor. This regime is currently out of range of experimentation using accelerators, but such phenomena might be manifested in the dense cores of neutron stars and, therefore, might be open to study through astro-

nomical observation.

The most useful experimental approach in the area of high temperature QCD matter is the detailed analysis of heavy ion collisions. The value of ~ 1 GeV/fm³ is not much higher than the energy density inside nucleons (~ 500 MeV/fm³) and nuclei (~ 150 MeV/fm³), and it is also comparable to estimates of the initial energy density created in hadronic collisions at high energy accelerators. In heavy ion collisions at relativistic velocities there is both a compression of the baryonic matter in the nuclei and a release of a large amount of energy within a small volume from the almost simultaneous collision of many nucleons. One or the other, or both, of these consequences of the interactions have the potential to produce new forms or phases of QCD matter. This is one of the prime reasons why in the past few decades much effort has been spent studying collisions of heavy ions at higher and higher energies [15]. Another important aspect of such studies is the extraction of valuable information about the mechanisms of particle production in small and large systems at high energies.

The most recent experimental facility for the study of heavy ion collisions is the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory. Since the inception of the physics program in July, 2000, four experiments at RHIC, namely BRAHMS, PHENIX, PHOBOS, and STAR, have studied the collisions of p+p, d+Au, and Au+Au at center-of-mass collision energies per incident nucleon pair, $\sqrt{s_{NN}}$, from 19.6 to 200 GeV. Data from all four detectors are being studied to get a better understanding of the physics of heavy ion collisions, and, in particular, to search for evidence of the creation of new forms of QCD matter [16]. To the best of our knowledge, where there is overlap, there are no major differences in the data and extracted results obtained by the four experiments at RHIC. The level of agreement is a testament to the quality of the detectors and the analyses performed by the collaborations and is a great strength of the whole RHIC research program. This paper summarizes the most important results obtained to date by the PHOBOS collaboration and the conclusions that can be drawn from PHOBOS results, augmented where necessary by data from other experiments.

One of the most important results found at RHIC is the evidence that in central Au+Au collisions at ultra-relativistic energies, a strongly interacting medium is created with extremely high energy density, whose description in terms of simple hadronic degrees of freedom is inappropriate. These conclusions are based on very general and, to a large extent, model independent arguments. From the smooth evolution of all observables studied up till now as a function of energy, system size, and rapidity, it appears likely that the transition to this new state is a cross-over.

It is not claimed that the observed phenomena are unique to RHIC energies. Nor is it claimed that to date there is direct evidence in the data for color deconfinement or chiral symmetry restoration. It should be noted

that interpretations of the data which invoke a high density of gluons or other non-hadronic components are certainly consistent with, and could be construed to provide at least circumstantial evidence for, deconfinement. Also, the definition of the concept of deconfinement is not so clear in a strongly interacting medium. No convincing evidence has been found for the creation at RHIC of the wQGP, as naïvely imagined by a large part of the heavy ion community in the era before the start of the RHIC physics program. It should be noted that this expectation may have resulted from a misinterpretation of the lattice results. The calculations reveal that the pressure and energy density reach 70–80% of the Stefan-Boltzmann value (i.e. the value for a non-interacting gas) for temperatures above the critical temperature (see, as one recent example, [17]). This observation was typically assumed to imply the presence of a weakly interacting system although questions were occasionally raised (for one early example, see [18]). More recently, this conclusion has been seriously challenged (see, for example, [19]). It should be noted that this recent reversal of opinion was to a large degree driven by the experimental results from RHIC. Recent lattice QCD studies have shown that the quarks do retain a degree of correlation above the critical temperature (see, for example, [20]). However, at extremely high energy density (for example, the very early universe), the theoretical expectation remains that the system will become weakly interacting.

Another equally interesting result from RHIC arose from the studies of the mechanism of particle production in nuclear collisions. Specifically, it has been discovered that much of the data in this new regime can be expressed in terms of simple scaling behaviors. Some of these behaviors had been noted in data at lower energies or for simpler systems. These observations suggest either the existence of strong global constraints or some kind of universality in the mechanism of the production of hadrons in high energy collisions, possibly connected to ideas of parton saturation. The data strongly suggest that the initial geometry and very early evolution of the system establish conditions which determine the final values of many observables. The most concise formulation of this discovery is the statement that the overall properties of the data appear to be much simpler than any of the models invoked to explain them. A full exploration and detailed analysis of all aspects of the data will be required for a complete understanding of the properties of QCD physics in the interesting regime probed by heavy ion collisions at relativistic velocities.

Section II of this paper describes the derived properties of the state formed shortly after the collisions at RHIC, Sect. III describes the evidence that this state is strongly interacting, and Sect. IV discusses the broad range of scaling behaviors that have been discovered.

As a useful reference, the PHOBOS detector and its properties are briefly described in Appendix A. Variables used in the description of the data, in particular those relating to event characterization, are defined in Appendix

B. The precise determination of the collision impact parameter or centrality is critical to heavy ion physics in general and the PHOBOS program in particular. Appendix C describes how centrality and the biases associated with triggering and various elements of the data analysis are derived from measurements and simulations for the various colliding systems and beam energies.

II. PROPERTIES OF THE INITIAL STATE PRODUCED AT RHIC

The primary goal of the RHIC accelerator was the study of QCD matter under extreme conditions. In particular, it was hoped that the center-of-mass energies more than an order of magnitude higher than achieved at the SPS accelerator at CERN would lead to the creation of a system with significantly higher energy density. An additional consequence of the higher beam energy compared to the SPS was the displacement of the projectile baryons a factor of two farther apart in rapidity. This was expected to lead to a lower baryon chemical potential in the high energy density region at mid-rapidity. Although progress has been made recently in lattice calculations which include quarks (see, as one example, [21]), the most extensively studied system remains one with a baryon chemical potential close to zero (see, for example, [22] and references therein). Therefore, creation of a system with a lower baryon chemical potential might offer the potential for more reliable comparisons of experimental data to the fundamental QCD predictions. This section describes the conclusions that can be drawn from PHOBOS data concerning these two critical properties of the state formed in collisions of heavy ions at RHIC.

A. Energy Density

In very high energy heavy ion interactions, the maximum energy density occurs just as the two highly Lorentz contracted nuclei collide. Clearly this system is very far from being equilibrated and, as a result, the value of the energy density, although well defined, may not be very interesting. In any reference frame, the potentially more interesting quantity is the energy density carried by particles which are closer to equilibrium conditions, i.e. those particles which have, on average, comparable longitudinal and transverse momenta. These conditions are roughly equivalent to restricting the particles to a range of pseudorapidity $|\eta| \leq 1$. Unfortunately, there are no direct measures of energy density and, therefore, it must be inferred from the properties of the detected particles. PHOBOS data have been used to investigate what range of initial energy densities are consistent with the observations. Studies of pseudorapidity and transverse momentum distributions, as well as elliptic flow, have been combined to constrain assumptions about the energy in the system and the time evolution of the volume

from which the particles emanate.




FIG. 1: Pseudorapidity density of charged particles emitted in Au+Au collisions at three different values of the nucleon-nucleon center-of-mass energy. Data are shown for a range of centralities, labeled by the fraction of the total inelastic cross section in each bin, with smaller numbers being more central. Grey bands shown for selected centrality bins indicate the typical systematic uncertainties (90% C.L.).

Figure 1 shows distributions of charged particle pseudorapidity densities, $dN_{ch}/d\eta$, for Au+Au collisions at $\sqrt{s_{NN}}=19.6, 130$, and 200 GeV for various centralities [24]. The produced particle densities are at their maximum near mid-rapidity and increase with both collision energy and centrality. The right panel of Fig. 2 is a compilation of the evolution of the mid-rapidity particle density $dN_{ch}/d\eta|_{|\eta|\leq 1}$, per participating nucleon pair,

$N_{part}/2$, as a function of collision energy [25, 26]. It should be noted that multiplicity measurements were initially obtained by PHOBOS and later confirmed by the other experiments at every new beam energy and species provided during the first three RHIC runs, from the first Au+Au collisions [27] through the d+Au collisions [28].

It is interesting to note that the measured mid-rapidity charged particle density at RHIC is lower than the prediction of most models (see the left panel of Fig. 2 as well as [23, 29] and references therein). Two of the models which predicted a value close to that seen in the data invoked the concept of parton saturation in the initial state [30, 31]. These models were similar to more recent formulations which describe the formation of a Color Glass Condensate (CGC). This newer CGC model successfully related the pseudorapidity and energy dependences of charged particle production to the gluon structure function measured in e+p collisions [32]. It should be noted that this model also made predictions concerning the production of high transverse momentum particles in d+Au collisions at mid-rapidity and forward rapidity [33–35] with mixed success when compared to the d+Au data [36–40]. The search for other evidence for possible parton saturation effects remains a topic of interest at RHIC but a more detailed discussion is beyond the scope of this paper.

Before attempting to make detailed estimates of the energy density, it is important to stress that the mid-rapidity particle density at the top RHIC energy is about a factor of two higher than the maximum value seen at the SPS [25] and there is evidence that the energy per particle has not decreased [41, 42]. Thus, with little or no model dependence, it can be inferred that the energy density has increased by at least a factor of two from $\sqrt{s_{NN}}=17$ to 200 GeV.

In addition to the measured particle multiplicities, estimating the energy density more precisely requires knowledge of the average energy per particle as well as the volume from which they originate. PHOBOS data for the transverse momentum distribution of charged particles [43] can be used to find a mean transverse momentum but these data only extend down to a few hundred MeV/c. Alternatively, Fig. 3 compares identified particle yields at very low transverse momentum measured by PHOBOS [44] to PHENIX data [45] for higher momenta. Both data sets are for particles emitted near mid-rapidity in central Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV. The PHOBOS data clearly demonstrate that the fits shown hold over the full phase space and that extrapolating over the full range of transverse momentum should give a correct value for the average. Accounting for the yields of the various particles, an average transverse momentum for all charged particles of $\langle p_T \rangle \sim 500$ MeV/c can be derived. The value found from the PHOBOS charged particle distributions is the same within 5%. Averaging over the pions, kaons, and nucleons, and assuming the yields for the unobserved neutral particles, an average transverse mass, m_T , of ~ 570 MeV/c² can be extracted. Under the




FIG. 2: (Left Panel) The PHOBOS measurement of charged particle density near mid-rapidity in central Au+Au at $\sqrt{s_{NN}}=200$ GeV (shown by the vertical line with the dashed lines denoting the systematic uncertainty) is compared to theoretical predictions. This panel is adopted from [23] which contains the full explanation and references. Following the example of [23], the PHOBOS $dN_{ch}/d\eta$ has been multiplied by 1.1 to convert to dN_{ch}/dy for comparison to theory. (Right Panel) Normalized pseudorapidity density of charged particles emitted within $|\eta| \leq 1$ in central Au+Au (AGS and RHIC) and Pb+Pb (SPS) collisions as a function of nucleon-nucleon center-of-mass energy [25, 26]. References for the lower energy data can be found in [25]. The densities have been divided by the average number of pairs of participating nucleons in the interactions. The data follow a simple logarithmic extrapolation from lower energies as shown by the line drawn to guide the eye.

assumption of comparable average transverse and longitudinal momenta, the average energy per particle is equal to the transverse energy at mid-rapidity. Since there are significant uncertainties in other elements of the calculation, a rounded estimate of 600 MeV per particle will be used.

The total energy in the system created near mid-rapidity in central Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV can be found from

$$E_{tot} = 2E_{part}dN_{ch}/d\eta|_{|\eta| \leq 1}f_{neut}f_{4\pi},$$

where E_{part} is the energy per particle, $dN_{ch}/d\eta|_{|\eta| \leq 1} = 650 \pm 35$ (syst) is the mid-rapidity charged particle density for the 6% most central collisions, f_{neut} is a factor of 1.5 to roughly account for undetected neutral particles, and the factor of 2 integrates over $-1 \leq \eta \leq +1$. One further issue to consider is that there are particles with similar total momentum in the center-of-mass system but which are not traveling predominantly in the transverse direction. The correction for these additional particles, $f_{4\pi}$, is trivially estimated from the fraction of phase space outside $\theta = 45^\circ - 135^\circ$ (i.e. outside $|\eta| \leq 1$) and equals about 1.3. Combining all these terms, the total energy contained in all particles emitted near

mid-rapidity, with transverse and longitudinal momenta consistent with emission from an equilibrated source, is about 1600 GeV. This is roughly 4% of the total energy of 39.4 TeV in the colliding system.

Converting this to a density in the rest frame of the system consisting of these particles requires knowledge of the volume within which this energy is contained at the earliest time of approximate equilibration. For central collisions, a transverse area equal to that of the Au nuclei (≈ 150 fm²) can be assumed, but which value to use for the longitudinal extent is not as clear. One extreme is to take the very first instant when the two Lorentz contracted nuclei overlap (longitudinal size ≈ 0.1 fm), which yields an upper limit on the energy density in excess of 100 GeV/fm³. There is, however, no reason to assume that at such an early instant the system is in any way close to equilibrium. A second commonly-used assumption is that proposed by Bjorken [46], namely a transverse size equal to the colliding nuclei and a longitudinal size of 2 fm (corresponding to a time of the order of $\tau \sim 1$ fm/c since the collision) which implies an energy density of about 5 GeV/fm³. Finally, the elliptic flow results discussed below suggest that an upper limit of the time for the system to reach approximate equilibrium is




FIG. 3: Transverse momentum distributions of identified charged particles emitted near mid-rapidity in central Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV. Invariant yield data shown are from PHENIX at higher momenta and PHOBOS at lower momenta. Boxes around the PHOBOS data indicate systematic uncertainties. Fits to PHENIX measurements are shown by solid curves ($\propto 1/[e^{(m_T/T)} + \epsilon]$, where $\epsilon = -1$ and $+1$ for mesons and baryons, respectively, m_T is the transverse mass, and T is the fit parameter). Note that the extrapolations (dashed curves) of the fit to the data at higher momenta are consistent with the low momentum yields.

of the order of 1–2 fm/c. Using the upper range of this estimate and further conservatively assuming that the system expands during this time in both the longitudinal and transverse directions, one obtains a lower limit of the energy density produced when the system reaches approximate equilibrium at RHIC of ≥ 3 GeV/fm³. Even this very conservative estimate is about six times the energy density inside nucleons and about twenty times the energy density of nuclei. Therefore, this is a system whose description in terms of simple hadronic degrees of freedom is inappropriate.

B. Baryon Chemical Potential

Turning to the baryon chemical potential, μ_B , early results regarding this property of the high energy density medium produced at RHIC came from the measurement of the ratios of charged antiparticles to particles near mid-rapidity for central collisions. In the simplest Boltzmann approximation, the ratio of antiprotons to protons is proportional to $e^{-2\mu_B/T}$, where T is the temperature at the time of chemical freezeout. Figure 4 compares the antiparticle to particle ratios for both protons and kaons to the corresponding numbers found at lower energies [47]. Clearly, the systems formed at RHIC are much closer to having equal numbers of particles and antiparticles than was true at lower energies. The measured value of $0.73 \pm 0.02(\text{stat}) \pm 0.03(\text{syst})$ for




FIG. 4: Ratios of identified antiparticles over particles measured near mid-rapidity in central collisions of Au+Au (AGS and RHIC) and Pb+Pb (SPS) as a function of nucleon-nucleon center-of-mass energy. Error bars are statistical only.

the antiproton to proton ratio near mid-rapidity for central Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV [48] indicates that these collisions are approaching a very low value of μ_B . Within the framework of thermal models, these ratios can be used to extract the baryon chemical potential [49]. Assuming a hadronization temperature of 165 MeV, a value of $\mu_B=27$ MeV was found for central Au+Au at $\sqrt{s_{NN}}=200$ GeV. This baryon chemical potential is an order of magnitude lower than what was obtained for Pb+Pb data at $\sqrt{s_{NN}}=17.2$ GeV from the SPS (see references in [47, 48]). Although the system created near mid-rapidity at RHIC cannot be described as completely baryon-free, it is clearly approaching the environment treated in most lattice calculations.

1. Comparison of particle ratios in Au+Au and d+Au

In addition to the higher center-of-mass energies, a critical element of the design of RHIC was the ability to collide asymmetric systems. This capability was first exploited with the collision of deuterons with gold nuclei at $\sqrt{s_{NN}}=200$ GeV. It is hoped that analysis of such simpler systems will serve as critical “control” experiments to aid in the understanding of the more complicated nucleus-nucleus data. As a first example, this section presents a study of the antiparticle to particle ratios.

As described above, particle ratios can be used to extract information about the properties of the system, in particular the chemical potentials. The measured values of these parameters are established at the point of chemical freeze-out when inelastic interactions between the produced particles cease. However, the properties of the early evolution of the system can clearly influence final conditions. Of particular interest in this regard is the ratio of antiprotons to protons measured at mid-rapidity. This ratio can be naively interpreted as reflecting the interplay of two mechanisms, namely the transport of




FIG. 5: The ratio of antiprotons to protons emitted in a rapidity region spanning approximately $0.0 < y < 0.8$ (where positive rapidity is in the direction of the deuteron projectile) for d+Au collisions at $\sqrt{s_{NN}}=200$ GeV. Data are shown for 4 centrality ranges. The parameter ν is the average number of collisions suffered by each participant in the deuteron (N_{coll}/N_{part}^d). Statistical and point-to-point systematic uncertainties are shown as bars and brackets, respectively. The results of several models are shown for comparison.

baryons from the two projectile nuclei to mid-rapidity and the production of antibaryon-baryon pairs in the interaction. By studying ratios as a function of centrality in d+Au, the effect of multiple collisions of the nucleons in the deuteron can be explored. The surprising result is shown in Fig. 5 [50].

The naïve expectation, supported by various model calculations, was that the proportion of antiprotons near mid-rapidity would fall slowly with collision centrality as the deuteron participants suffered more collisions and consequently were effectively transported closer to the center-of-mass rapidity. In contrast, the data show a ratio which is consistent with being the same at all centralities. At present, no simple explanation or interpretation of the observed particle ratios is known.

The d+Au data at RHIC serve an important function as a control experiment since it is generally believed that an extended volume of high density matter is not formed in these collisions. Understanding the basic mechanisms of baryon transport and baryon pair production will clearly be critical to a full description of heavy ion interactions.

C. Nature of the transition to the high density regime

The transition to the high density state at RHIC has not been observed to create abrupt changes in any observable studied to date, including, among others, charged particle multiplicity, elliptic flow, HBT, as well as derived quantities such as energy density and freeze-out parameters. This lack of a dramatic change in char-

acter may make it more difficult to delineate the exact boundaries of the onset of significant influence from non-hadronic degrees of freedom. However, it should be noted that this observation may be consistent with the expectations concerning the nature of the phase transition from the most recent lattice QCD calculations [17, 51, 52] which predict a rapid cross-over in the region of the phase diagram believed to be relevant for the systems created near mid-rapidity at RHIC.

It should be noted that indications of possible non-monotonic behavior in the energy evolution of some quantities were reported in the range $\sqrt{s_{NN}}=5-10$ GeV at the CERN SPS (see, for example, [53] and references therein). The extracted properties of the environment created near mid-rapidity in these lower energy collisions are significantly different from those found near mid-rapidity at RHIC, with energy densities at least a factor of 3–4 smaller and baryon chemical potentials an order of magnitude or more larger. A discussion of these results at lower energy falls outside the scope of this paper but future work in this area might prove important to the full exploration of the QCD phase diagram.

III. STRENGTH OF INTERACTIONS IN THE HIGH ENERGY DENSITY MEDIUM

In early discussions of the high density systems hoped to be formed in RHIC collisions, the expectation was that a deconfined state of quarks and gluons would be weakly interacting. This interpretation arose at least partly from the naïve assumption that any matter that attained a large fraction of the Stefan-Boltzmann limit for the pressure would act like a gas [19]. One of the most dramatic early discoveries at RHIC is the clear indication that the nature of the systems formed is very far from weakly interacting. Evidence for this conclusion is found in the magnitude of elliptic flow and in the centrality dependence of particle production at high transverse momentum. The former provides information on the manner in which particle production depends on the shape of the incident system and the latter explores how the spectrum of the produced particles is impacted by the medium. Additional evidence is provided by the strength of the radial flow signals seen in the yields of particles at very low transverse momentum, a measurement unique to PHOBOS.

Figure 6 shows PHOBOS measurements of the magnitude of elliptic flow, v_2 , near mid-rapidity ($|\eta| \leq 1$) in Au+Au collisions at $\sqrt{s_{NN}}=130$ and 200 GeV as a function of centrality, denoted by $\langle N_{part} \rangle$ [54, 55]. Two different methods of determining the flow signal, one based on counting hits in the multiplicity detector and one based on counting tracks in the spectrometer [55], were used at the higher beam energy. Similar results were first shown for RHIC data by the STAR collaboration [56]. Figure 7 shows data from the track-based method in the rapidity interval $0 < \eta < 1.5$ for the 50%




FIG. 6: Elliptic flow of charged particles near mid-rapidity ($|\eta| < 1$) as a function of centrality in Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV using two different methods (closed circles and triangles, see text for details) and at $\sqrt{s_{NN}}=130$ GeV (open triangles). Grey boxes show the systematic errors (90% C.L.) for the 200 GeV data. The curve shows the prediction from a relativistic hydrodynamics calculation.




FIG. 7: Elliptic flow of charged particles emitted near mid-rapidity ($0 < \eta < 1.5$) in the 50% most central Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV as a function of transverse momentum. Grey boxes show the systematic uncertainties of the data (90% C.L.). The curve is the prediction of a relativistic hydrodynamics calculation.

most central Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV as a function of transverse momentum, p_T [55]. Data in both figures are compared to the predictions of a hydrodynamical calculation [57]. These results show that elliptic flow is unexpectedly strong at RHIC energies. Over a wide range of centrality and transverse momentum, the value near mid-rapidity is as large as that calculated under the assumption that a boost-invariant relativistic hydrodynamic fluid was formed.

When two nuclei collide with non-zero impact parameter, the lenticular (or almond-shaped) overlap region has an azimuthal spatial asymmetry (see right panel of Fig. 35). However, if the particles do not interact af-

ter their initial production (presumably with azimuthally uniform momenta), the asymmetrical shape of the source region will have no impact on the azimuthal distribution of detected particles. Therefore, observation of azimuthal asymmetry in the outgoing particles is direct evidence of interactions between the produced particles. In addition, the interactions must have occurred at relatively early times, since expansion of the source, even if uniform, will gradually erase the magnitude of the spatial asymmetry.

Qualitatively, it is clear that an asymmetric system of interacting particles will have asymmetric pressure gradients which drive any observed particle asymmetry. Hydrodynamical models can be used to calculate a quantitative relationship between a specific initial source shape and the distribution of emitted particles (see, for example, [57]). Due to the ideal nature of the fluid assumed in these models, the resulting final asymmetry is generally believed to be an upper limit for a specific starting condition. From the strength of the observed elliptic flow and from the known dimensions of the overlap region in Au+Au collisions, it can be conservatively estimated that the pressure build-up in the initially formed medium must have occurred in a time less than about 2 fm/c [58]. Thus, the presence of a strong flow signal carries several important implications, the first of which, a limit on the timescale for equilibration, has been used previously in the discussion of energy density. In addition, one can conclude that at these early times the initially produced particles must have interacted strongly, corresponding more closely to the conditions in a fluid rather than a gas.

Additional indirect evidence that a strongly interacting system is produced in heavy ion collisions at RHIC is the observed yield of particles with very small transverse momentum (≤ 100 MeV/c) [44], shown previously in Fig. 3. Recall that the production of particles with p_T as low as 30 MeV/c were consistent with extrapolations from a fit to the distribution in the range of a few hundred MeV/c to a few GeV/c. If, in RHIC collisions, a medium of weakly interacting particles was initially produced, one could expect an enhancement in the production of particles with wavelengths up to the overall size of the collision volume (i.e. coherent pion production) [59]. In essence, the observation of no such excess is another manifestation of the high pressure gradient and strongly interacting nature of the medium, which gives rise to the large magnitude of the elliptic flow signal seen at RHIC. These properties would also produce strong radial flow so that any particles initially produced with low velocity would subsequently be accelerated by the interactions.

The study of the yield of particles with large transverse momentum can be used to more directly explore the level of interactions present in the medium produced in $\sqrt{s_{NN}}=200$ GeV Au+Au collisions at RHIC. Presuming that high momentum transfer processes are induced via relatively short-range interactions, one may expect QCD factorization theorems to hold and, therefore, a particular hard process can be induced by any binary collision

in the overall nucleus-nucleus interaction [60, 61]. This is the motivation for the nuclear modification factor, R_{AA} defined in Appendix B3 and first studied at RHIC by PHENIX [62, 63], which measures how effective each particular binary collision is for inducing a hard scattering process. Strong deviations from unity indicate violations of factorization, which may be caused by initial or final state effects. In their pioneering work, the PHENIX collaboration showed that in central collisions of Au+Au at $\sqrt{s_{NN}}=130$ GeV there was significant suppression of the yield of high transverse momentum particles compared to the p+p data scaled by the number of binary collisions, N_{coll} .




FIG. 8: Nuclear modification factor, R_{AA} , as a function of transverse momentum for Au+Au collisions at $\sqrt{s_{NN}}=62.4$ GeV (closed symbols) and 200 GeV (open symbols), for six centrality ranges. Centrality is expressed as a fraction of the total inelastic cross section with smaller numbers being more central. Bars and brackets show statistical and systematic uncertainties, respectively. The solid (dashed) line shows the expectation for scaling with N_{coll} ($N_{part}/2$) times p+p data.

The PHOBOS collaboration has confirmed that a similar effect is present in Au+Au collisions at 200 GeV [43], and has also performed the first similar studies at $\sqrt{s_{NN}}=62.4$ GeV [64], see Fig. 8. More importantly, as discussed later and shown in Fig. 31, the yields for Au+Au interactions at $\sqrt{s_{NN}}=200$ GeV which span a range of more than a factor of five in the number of participants were found to scale with the number of participants, when compared with central Au+Au collisions, to

within $\leq 25\%$ at all transverse momenta. The fact that data up to p_T of 4 GeV/c show much the same scaling as at low momentum clearly demonstrates that any scaling of hard processes with the number of binary collisions is completely obliterated. Note the significant difference in the shapes of R_{AA} at the two energies shown in Fig. 8, as well as the fact that the difference is similar at all centralities. Additional discussion of this interesting observation as well as other scaling properties of the data can be found in Sect. IV.

It is important to note that, except where specifically mentioned, the reference p+p data in this and all other cases of comparison to RHIC data is for inelastic collisions. This choice is made for consistency rather than being strongly motivated by physics considerations. In most cases, the difference between the yield in non-single diffractive (NSD) and inelastic measurements is about 10% or less.




FIG. 9: Nuclear modification factor, R_{dAu} , as a function of transverse momentum for d+Au collisions at $\sqrt{s_{NN}}=200$ GeV, for four centrality ranges. Centrality is expressed as a fraction of the total inelastic cross section with smaller numbers being more central. Bars and brackets show statistical and systematic uncertainties, respectively. The shaded area shows the uncertainty (90% C.L.) in R_{dAu} due to the systematic uncertainty in N_{coll} and the scale uncertainty in the proton-proton data. In the bottom right panel, the nuclear modification factor, R_{AA} , for the 6% most central Au+Au collisions at the same energy is shown as a dark curve for comparison.

As mentioned above, the observed suppression of hard processes could result from some modification in the initial state (see, for example, [33]) as well as from interactions in the dense medium formed after the collision. To investigate this possibility, similar data were taken for d+Au collisions at the same energy. Figure 9 shows the nuclear modification factor, R_{dAu} , measured

by PHOBOS in d+Au at $\sqrt{s_{NN}}=200$ GeV, in four different impact-parameter ranges and the similar modification factor, R_{AA} , in central Au+Au collisions at the same energy [36]. Note the dramatic difference between the results for central d+Au and Au+Au collisions at higher transverse momentum shown in the lower right panel of the figure. For $2 \text{ GeV}/c \leq p_T \leq 6 \text{ GeV}/c$ the yield of charged particles in d+Au is consistent with binary collision scaling of p+p data, whereas in Au+Au collisions the yield is clearly suppressed.

The interpretation of the observed binary collision scaling at higher p_T in d+Au (i.e. the fact that the data points at higher p_T in Fig. 9 are similar at all centralities and all lie near unity) is unclear since it might be a consequence of the interplay of an enhancement (similar to the so-called ‘‘Cronin effect’’), and some suppression, due to either energy loss in the final state or parton saturation effects in the initial state. Furthermore, it should be noted that the number of participants and the number of collisions do not deviate as much with centrality in d+Au as in Au+Au. Using the number of participant pairs as the scaling variable (i.e. using $R_{dAu}^{N_{part}}$ defined in Appendix B3) would raise the values at all transverse momenta by an average factor of about 1.65. However, the factor would differ only by 29, 14, and 6% for centrality bins of 70–100, 40–70, and 20–40%, respectively, compared to the 0–20% data. These shifts are comparable to, or smaller than, the systematic uncertainties in the overall scale of the modification factors. Thus, the observation of similar levels at all centralities does not necessarily imply scaling with the number of collisions.

To further complicate the interpretation, the value of the nuclear modification factor was found to depend on the pseudorapidity of the emitted particles. This was first inferred from the comparison of the PHOBOS results to those of the other experiments. It can also be seen from PHOBOS results directly as shown in Fig. 10 [37]. Data from BRAHMS suggest that this trend may continue to higher positive pseudorapidity [38] while preliminary PHENIX data suggests that R_{dAu} may even continue rising for negative pseudorapidity (i.e. towards the Au projectile rapidity) [39]. The trend seen in the PHOBOS and BRAHMS data has been interpreted as support for the CGC model, but this conclusion is far from clear and the PHENIX data at negative pseudorapidity remain even more mysterious [40]. For this reason, the observation of the particular value of $R_{dAu}=1$ at higher p_T is an accident of the PHOBOS acceptance and again does not necessarily imply N_{coll} scaling.

Therefore, the important feature is not the possible scaling of the particle yields in d+Au with N_{coll} times p+p yields, but instead the very significant difference between the transverse momentum dependence of the d+Au and Au+Au nuclear modification factors. The larger system appears to lead to a strong suppression while the smaller system does not. Very similar results were reported simultaneously by all four RHIC experiments [36, 65–67]. Part of the difference in the behavior




FIG. 10: Nuclear modification factor, R_{dAu} , for four different values of p_T as a function of pseudorapidity in d+Au collisions at $\sqrt{s_{NN}}=200$ GeV. PHOBOS results away from mid-rapidity [37] are compared to data near $\eta=0$ from BRAHMS [65], PHENIX [66], and STAR [67]. For the PHOBOS points, the error bars are the point-to-point systematic errors. The systematic errors in the overall scale of the PHOBOS R_{dAu} are shown as grey bands.

of the two colliding systems may be attributed to initial state effects. However, it is difficult to avoid the conclusion that the majority of the difference in R_{AA} compared to R_{dAu} results from the impact of the high energy density matter on the yield of particles with p_T in this measured range. Clearly, the constituents of the medium produced in the central Au+Au collisions experience a significant level of interaction. Since, as discussed above, the system at this early stage cannot be primarily hadronic in nature, one can conclude that the high energy density matter created at RHIC interacts very strongly with high p_T partons (or with whatever constituents comprise the dominant degrees of freedom at this early stage). It certainly does not appear to be a weakly interacting parton or hadron gas.

In related measurements, the STAR experiment has studied back to back correlations of high p_T particles. Measuring the yield of particles as a function of the azimuthal angle relative to a very high p_T trigger particle, a suppression was found in particles emitted on the opposite side [68]. This suppression was found to depend on the azimuthal angle of the trigger particles with respect to the reaction plane [69]. One strength of the correlation analysis is that it is essentially self-normalizing in the sense that the result does not depend on any assumptions about the scaling of the primary production of particles. One can interpret this as additional support for the conclusions that are being drawn from the single particle data.

Further evidence that the system may be both non-hadronic in nature and also strongly interacting comes from flow data for identified particles. PHENIX [70] and STAR [71] have measured the elliptic flow and its dependence on transverse momentum for a variety of mesons and baryons. These data appear to be consistent with an interpretation that the flow of produced particles results from the recombination of quarks which are themselves flowing [72]. The impact of this flow of quarks is that the v_2 parameter divided by the number of valence quarks scales as a function of the transverse momentum, also divided by the number of valence quarks. It should be noted that this recombination model only holds for elliptic flow at higher values of $p_T \geq 1 - 2$ GeV/c. If this interpretation is correct, it lends support to the presumption that the system has a component of constituent quarks which experience significant interactions early in the evolution of the collision.

In conclusion, the data from RHIC collisions provide strong evidence for the creation of a very high energy density, relatively baryon-free, medium which cannot simply be described in terms of hadrons and whose constituents experience significant interactions with each other.

IV. SIMPLE SCALING BEHAVIORS OF PARTICLE PRODUCTION

The wide range of systems and energies provided by the RHIC accelerator, combined with the unique capabilities of the PHOBOS detector, has allowed a study of the properties of particle production over a very broad range of initial conditions and particle phase space. In the process of this study, a surprising result was discovered. It emerged that an enormous span of data for charged particles emitted in d+Au and Au+Au collisions at RHIC energies could, to a large extent, be described using only a few simple unifying features. Some of these scaling behaviors had been observed previously, either at lower energies or for less complicated systems than heavy ion collisions. This section describes the extent to which these scaling behaviors and other unifying features apply to charged particle production at RHIC energies.

In a wide variety of systems (hadron+A up to A+A), the total number of emitted charged particles is observed to have a very simple dependence on energy and centrality. In all cases, the total multiplicity appears to scale linearly with the number of participant pairs, $N_{part}/2$. It should be noted that throughout this document the generic term “participant pairs” refers simply to the total number of participants divided by 2, i.e. a quantity that is unity in p+p, and does not imply a matched pair from the two colliding species. The total multiplicity of charged particles emitted in hadron+A (including p+A) and d+A is equal to $N_{part}/2$ times the multiplicity observed in p+p. In contrast, for heavier nucleus-nucleus interactions, the constant of proportionality is the multiplicity produced in e^+e^- annihilations, which is ap-

proximately equal to that measured in p+p at twice the center-of-mass energy. This is suggestive of a universal energy dependence of charged particle multiplicities in strong interactions. Centrality, as reflected by the number of participants (both the total number and, for asymmetric systems, the number in each of the nuclei) appears to have a strong influence on the shape of the pseudorapidity distributions. In addition, the yield of particles produced in A+A collisions at a variety of centralities, when compared to the yield in central collisions, displays a dependence on the number of participants out to a surprisingly large value of transverse momentum.

Over a broad range of emission angles, the distributions of pseudorapidity density and the elliptic flow signal, when measured as a function of the variable $\eta' = \eta - y_{beam}$ (i.e. when viewed in the approximate rest frame of one of the colliding particles), appear to be identical both in shape and magnitude at all beam energies over a large range of η' . The details of the shape of the distributions depend on the impact parameter, but again in an energy-independent way. In addition to this extensive longitudinal scaling, no evidence is seen for a boost invariant central plateau in the pseudorapidity distributions of either particle multiplicity or elliptic flow.

Finally, many properties of particle production exhibit separate dependences on the energy and centrality of the collisions which factorize to a surprising degree. In other words, the centrality dependence of data such as pseudorapidity density and transverse momentum spectra was found to be identical even at center-of-mass energies separated by up to an order of magnitude. Another aspect of the centrality dependence is the observation that many differences between data for Au+Au and p+p persist essentially unchanged over a centrality range corresponding to a number of participants that spans a factor of 5 or more.

A. Energy dependence of total multiplicity

The most basic observable in the study of multiplicity is the total number of produced particles. Collisions at RHIC extend the center-of-mass energy range available in heavy ion interactions by more than an order of magnitude. Section II A described the energy dependence of the mid-rapidity particle density. In this section, the total integrated particle yield is discussed. As is clearly shown in Fig. 1, the PHOBOS multiplicity detector extends over a uniquely broad range of pseudorapidity and therefore the extrapolation to account for missing regions of phase space is small even at the highest RHIC energy. The total multiplicity of charged particles per participant pair in A+A collisions over a wide range of energies is shown in Fig. 11, along with data from $p(\bar{p})+p$, d+Au and e^+e^- annihilation into hadrons [73]. The d+Au value has also been divided by the number of participant pairs. The nucleus-nucleus data are for central collisions. However, this choice is inconsequential since, as will be

discussed in the following section, the total multiplicity per participant pair appears to be approximately independent of centrality.

The various sets of data appear to have very different trends. The p+p (open squares) and d+Au (open star) data are consistently about 30% below the e^+e^- data, as shown in the lower panel, where all of the data points are divided by a fit to the e^+e^- data. Starting at the lowest energies, the A+A data rise much faster than both p+p and e^+e^- but then the slope of the energy dependence changes and above $\sqrt{s_{NN}} \sim 20\text{--}30$ GeV, the A+A data follow the trend of the e^+e^- data. The lower panel of the figure shows that these two sets agree to within 10% over a span of an order of magnitude in center-of-mass energy.

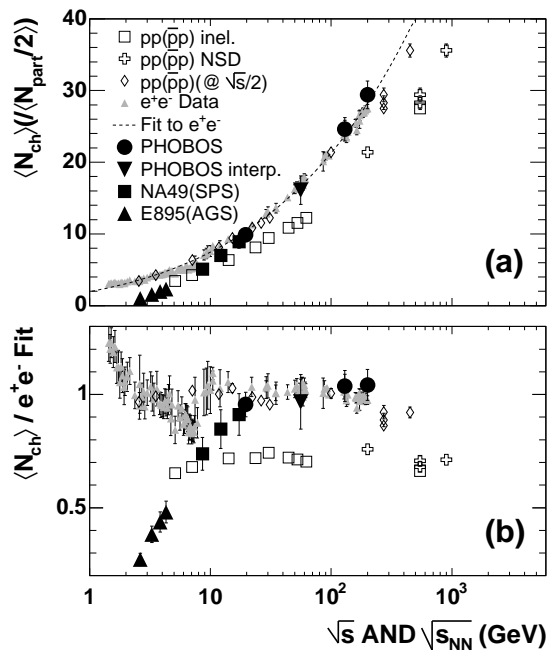


FIG. 11: Normalized total multiplicities of charged particles emitted in e^+e^- , p+p, d+Au, and Au+Au (AGS [74] and RHIC) and Pb+Pb (SPS [75]) collisions at a variety of nucleon-nucleon center-of-mass energies. Nucleus-nucleus data are all for central collisions and the multiplicities have been divided by the number of participating nucleon pairs. In the bottom panel, the values for all systems have been divided by a fit to the e^+e^- data. (Note: dAu point to be added)

The difference between the p+p and e^+e^- data can possibly be explained by accounting for the “leading particle effect”. The distribution of protons in x_F (See Appendix B 2 for definition) for p+p collisions at different energies was found to be approximately flat (with a spike at $x_F=1$ for elastic events) (a summary of these data can be found in [76]). One interpretation of these data is that a leading nucleon typically carries half of the beam energy. In p+p collisions, the x_F of the leading proton was found to directly anticorrelate with the particle mul-

tiplicity, as if the leading particle simply removed energy that would otherwise go into particle production [77–79]. By rescaling the center-of-mass energy for the p+p data by a factor of two (see open diamonds in Fig. 11), one observes that the multiplicities of p+p and e^+e^- reactions agree more closely over much of the energy range.

In contrast with the p+p data, which agree with the e^+e^- data over a large energy range only after rescaling, there is reasonable agreement of the total charged particle multiplicities between e^+e^- and A+A collisions over \sqrt{s} and $\sqrt{s_{NN}}$ of about 20 to 200 GeV with no rescaling. At lower energies, one sees an apparent “suppression” of the A+A multiplicity compared to both p+p and e^+e^- . This might be explained by reference to the substantial baryon excess found in the particle yields at these lower energies (e.g. the antiproton/proton ratio $\ll 1$) (see references in [48]). The relatively larger number of baryons compared to pions should tend to suppress the overall multiplicity since the baryon chemical potential reduces the entropy. Essentially, the conserved baryons themselves take up an increasing fraction of the available energy. Additionally, the overlap of the peak of the rapidity distributions of the net baryons and the produced pions [80] could result in increased pion absorption during the evolution of the system.

The arguments made here suggest that the total multiplicity per participant pair is a universal function of the available energy, irrespective of the colliding system [73]. This is a surprising result if p+p collisions are expected to be a “reference system” while the enhanced multiplicity in A+A is related to more exotic physics. Moreover, the prediction of the e^+e^- multiplicity is widely understood as a paradigmatic success of perturbative QCD, rather than resulting from the “soft physics” expected to dominate heavy ion collisions.

This interpretation of the comparison of p+p and Au+Au systems is validated by the $\sqrt{s_{NN}}=200$ GeV d+Au results from PHOBOS [81] shown in Fig. 11 for the most central collisions. If it takes more than one collision in order for all of the energy to be available for particle production, then one would expect the participants in the deuteron to contribute approximately half the multiplicity of an e^+e^- collision (i.e. with effective energy of \sqrt{s}), while the participants in the gold nucleus would contribute half a p+p collision. For a central d+Au collision, the ratio of gold to deuteron participants is approximately 8, so the “p+p-like” collisions should dominate, making the multiplicity closer to p+p, an expectation that is validated by the data.

It should be emphasized that this result applies mainly to the total multiplicity and not necessarily to other details of particle production. In other words, this argument does not imply that A+A collisions are merely scaled up e^+e^- annihilations. The presence of elliptic flow and strangeness enhancement, along with other observations, precludes this possibility. Furthermore, it is not argued that all observables in A+A collisions should be compared to similar data from p+p at twice

the center-of-mass energy. Still, the similarities between the total charged particle multiplicities of these various systems raise the question of what are the decisive differences between the larger and smaller systems. Some insight may come from studying the role of the size and shape of the collision volume, which will be addressed in later sections.

While the physics scenario as stated is consistent with a broad range of multiplicity data, it is complicated somewhat by the recent BRAHMS result on the net baryon distribution, which is interpreted in terms of the net rapidity loss of the incoming baryons [82]. Although the measurements do not include the bulk of the net baryons, the data can constrain the shape of the distributions substantially. The BRAHMS analysis finds that the average rapidity loss of the net baryons in central Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV is $\Delta y \sim 2$ units [82], which is consistent with values extracted from p+A data at lower energy [83, 84]. When translated into “available” energy, i.e. the total incoming energy minus the energy of the net outgoing baryons, only about 75% of the energy is left for particle production in central A+A collisions. It should be noted that this value is a lower limit based on the assumption that the effects of longitudinal expansion can be ignored. If this reduced available energy is accounted for in the Au+Au data as was done for p+p, the resulting data points in panel (b) of Fig. 11 would increase by about 15%. This would imply that Au+Au collisions are, in fact, able to convert the same amount of energy into a slightly larger number of particles than are produced in e^+e^- annihilations. Given the systematic uncertainties in the various data sets, it is difficult to determine which of these interpretations is correct. Furthermore, given the current lack of understanding of the longitudinal dynamics in RHIC collisions (see Sect. IV E), the validity of the assumption that all of the energy carried by the net baryons is “unavailable” for particle production is far from obvious. What is unambiguous is the surprisingly close correspondence of all systems despite the common assumption that somewhat different physics dominates in each case.

In summary, the data show that the systematics of the total charged-particle multiplicities are suggestive of a universal mechanism which affects “bulk” features of particle production in strongly-interacting systems. The dominant control variable in this picture appears to be the available or “effective” energy, per participant pair, which is apparently 50% of $\sqrt{s_{NN}}$ in a p+p or d+Au collision, but appears to be a significantly larger fraction of $\sqrt{s_{NN}}$ in A+A and presumably all of \sqrt{s} in e^+e^- reactions. This may simply be related to the fact that typical participants in an A+A collision are multiply struck when passing through the oncoming nucleus. A more complete description would involve a full explanation of the nature and origin of the outgoing baryons in both nucleon-nucleon and nucleus-nucleus collisions. All of these issues thus require a more comprehensive understanding of the early-time dynamics of the collision

process, including both the dynamics of baryon-number transport and entropy production.




FIG. 12: Total integrated charged particle multiplicity per participant pair as a function of number of participants. Data are shown for Au+Au collisions at $\sqrt{s_{NN}}$ of 19.6, 130 and 200 GeV [73] as well as d+Au [81] and $\bar{p}+p$ at 200 GeV [85]. The vertical bars include both statistical and systematic (90% C.L.) uncertainties.

B. Centrality dependence of total multiplicity

One of the key tools for understanding particle production in high energy p+A and A+A collisions is the study of the system-size dependence, either by varying the size of the colliding nuclei or by classifying the collisions according to centrality. Variation of the collision centrality not only changes the volume of the particle production region, but also the number of binary collisions per participant (see Appendix B 3 for more discussion of this topic). In addition to changing the collision energy, varying centrality therefore provides another handle, in principle, for changing the balance of particle production between ‘soft’ low-momentum processes and point-like ‘hard’ processes with large momentum transfer.

One of the more striking features of total particle production in Au+Au collisions at RHIC is the proportionality of the total charged-particle multiplicity to the number of participant pairs [73], as shown in Fig. 12 and compared to $\bar{p}+p$ and d+Au collisions. The figure also shows that the total charged particle multiplicity is proportional to the number of participating nucleons in Au+Au collisions at all three energies from $\sqrt{s_{NN}}=19.6$ to 200 GeV. The data suggest that the transition between p+p collisions and Au+Au is probably not controlled simply by the number of participants, as even very central d+Au collisions do not show any sign of trending up towards the level of the Au+Au data. As discussed in the preceeding section, this aspect of the total multiplicity is expected in the “available energy” *ansatz* since the Au participants, which dominate the total in d+Au, are expected to be more “p+p-like”.




FIG. 13: Ratios of total particle multiplicity data for a wide range of hadron-nucleus [86], and nucleus-nucleus collisions [26, 81] over the multiplicity in proton (or antiproton)-proton interactions [85, 87–89] are plotted versus the number of participating nucleons. The denominator for interactions induced by mesons, protons, or deuterons is proton-proton data at the same center-of-mass energy. For Au+Au interactions, the denominator is proton(or antiproton)-proton data at twice the center-of-mass energy. The error bars include both statistical and systematic effects. Furthermore, they are partially correlated due to common errors in N_{ch}^{pp} . Note that all the data fall on a common line with a slope of 1/2 (as expected since p+p has two participants) and zero intercept.

This topic represents one area where data for collisions of lighter nuclei at RHIC could make an important contribution. Extrapolation of Au+Au analysis to very peripheral collisions inevitably suffers from considerable systematic uncertainty in the number of participants. Lessons learned from analysis of lower energies and smaller systems such as d+Au are currently being applied in an attempt to reduce those uncertainties. However, it is clear that lighter systems, currently a high priority for future runs at RHIC, will provide vital input to the interpretation of these results.

Further information about the centrality dependence is shown in Fig. 13, the inset of which shows a detailed comparison of the PHOBOS d+Au results at $\sqrt{s_{NN}}=200$ GeV [81] with $\pi+A$, $K+A$, and $p+A$ for $\sqrt{s_{NN}} \approx 10$ –20 GeV [86]. In all cases in the inset, the total charged particle multiplicity in hadron-nucleus colli-

sions is divided by the p+p multiplicity at the same collision energy. Within the experimental uncertainty, scaling with the number of participant pairs times the data for p+p at the same energy is observed for all hadron-nucleus systems as was first recognized in earlier work [90, 91]. The range in N_{part} over which this scaling is shown to apply is extended significantly by the PHOBOS charged particle multiplicity in d+Au collisions versus centrality.

A similar analysis of Au+Au data for collisions at $\sqrt{s_{NN}}=19.6$ GeV and 200 GeV is shown in the main part of Fig. 13, again exhibiting scaling of the total multiplicity with the number of participant pairs, but in this case multiplied by p(p)p data at twice the center-of-mass energy. A particularly striking feature, as discussed in Section IV A, is the fact that, for all these systems and energies, the total number of charged particles is directly given by the number of participant pairs times the number seen in p+p after accounting correctly for the energy carried away by the leading baryon.

This continuation of the previously-observed approximate N_{part} scaling, which is now seen to apply to all systems and over an expanded range of energies from $\sqrt{s_{NN}}$ below 10 GeV to the highest at RHIC, represents one of the more surprising features of particle production at RHIC.

C. Centrality dependence of pseudorapidity distributions

It should be stressed that the universal N_{part} scaling of the total number of particles produced in Au+Au collisions does not result from rapidity distributions whose shape is independent of centrality, or N_{part} . The rapidity distributions do depend on both centrality and on the nature of the colliding systems, as is evident from Fig. 14 for Au+Au and Figs. 15 and 16 for d+Au. However, the dependence of the shape on centrality, as first reported in [92], is very specific.

The Au+Au pseudorapidity distributions shown in Fig. 14 appear to exhibit a sort of incompressibility in rapidity space. Thus a reduction in the number of particles at mid-rapidity is balanced by a similar increase of the number of particles at high rapidities, with the total number remaining constant. Obviously, moving particles around in rapidity changes the total longitudinal energy in the system. If the total energy available for produced particles depends only on the number of participants, energy must be conserved by changes in the distribution of transverse momentum.

The centrality dependence of pseudorapidity distributions in asymmetric systems can be studied using PHOBOS data for d+Au collisions as shown in Figs. 15 and 16. With increasing centrality, an increase in particle production (see Fig. 15) and a significant change in shape of the distributions (see Fig. 16) is observed. It should be stressed that the appearance of a “double-hump” structure in the d+Au distributions is primarily due to the




FIG. 14: Distributions of normalized pseudorapidity densities of charged particles emitted in Au+Au collisions at two energies and two ranges of centrality. The data have been divided by the average number of pairs of participating nucleons for each energy and centrality range. The centrality is designated by the fraction of the total inelastic cross section, with smaller numbers being more central. Only statistical errors are shown.

effect of the Jacobian associated with the transformation to $dN/d\eta$ from dN/dy (see related discussion in Section IV E 3). Although the shape changes in a non-trivial way, the integral of these distributions, when extrapolated to full phase-space, is found to be proportional to the number of participating nucleons, as was shown for many systems and energies in Section IV B.

The comparison of total particle multiplicity in d+Au and p+p can be extended by studying the ratio $dN/d\eta(\text{d+Au})/dN/d\eta(\text{p+p})$ as a function of pseudorapidity, as shown in Fig. 17. The main panel of the figure shows this ratio for various d+Au centralities, as a function of pseudorapidity. The inset and the arrows at the lower right demonstrate that, as was seen in p+A at lower energy [86, 91, 93–95], the data are consistent with a picture in which the density of produced particles which have a rapidity in the vicinity of the incident deuteron (gold) is proportional to the number of deuteron (gold) participants. The data suggest that the overall rapidity distribution, not just the integral of the distribution, is strongly influenced by the collision geometry.

In light of the discussion of particle production as a function of available energy in Sect. IV A, one might naïvely expect the ratio at positive rapidity in Fig. 17 to increase faster than the number of deuteron participants. This is because each deuteron participant interacts with multiple Au participants and is therefore “Au+Au-like” while each Au participant suffers far fewer collisions and is therefore “p+p-like”. Recall that the normalized multiplicity per participant pair in Au+Au collisions was higher than that in p+p collisions at the same center-




FIG. 15: Distributions of pseudorapidity densities of charged particles emitted in d+Au collisions at $\sqrt{s_{NN}}=200$ GeV for a variety of centralities. The positive pseudorapidity direction is that of the deuteron. The centrality is designated by the fraction of the total inelastic cross section, with smaller numbers being more central. Grey bands indicate the systematic uncertainties (90% C.L.).

of-mass energy. However, it is important to keep in mind that the detailed shape of the distribution, not just the relative height at the two ends, is a complicated function of centrality. For example, it has long been known that in p+A collisions, the yield of all particles with rapidity within a unit or so of that of the proton falls with increasing target mass [96]. Thus, one should not naïvely expect conclusions from integrated yields to apply simply to narrow fixed regions of pseudorapidity.

The longitudinal properties of particle production, and in particular the dependence on center-of-mass energy, are discussed in more detail in Sect. IV E.

D. Comparison of Au+Au and other systems

Figure 11 showed that the total charged particle multiplicities in the e^+e^- and A+A systems are very similar at a given center-of-mass energy, while those for p+p are somewhat smaller. To expand the comparison of these three very different systems, it is interesting to consider the full distributions of pseudorapidity. However, this study is complicated by the fact that the shapes of the Au+Au data vary dramatically with centrality (as is most clearly evident in Fig. 14). Figure 18 compares $dN_{ch}/d\eta$ normalized by the number of participant pairs for the 3% most central Au+Au collisions to inelastic




FIG. 16: The data of Fig. 15 is shown but in this case divided by the average number of participant pairs in each centrality bin. Systematic errors are not shown.

data for p+p and the distribution of dN_{ch}/dy_T (see definition in Appendix B 2) in the e^+e^- data [97], all at a $\sqrt{s_{NN}}$ or \sqrt{s} of 200 GeV [73]. The bottom panel of the figure demonstrates that the lower total multiplicity seen in p+p results from a pseudorapidity distribution that is suppressed by roughly a constant factor over all emission angles. The figure shows agreement in the overall rapidity distribution between A+A and e^+e^- . In comparing the two distributions, one should keep in mind the centrality dependence in the shape for Au+Au as well as the difference between dN/dy_T and $dN/d\eta_T$. Studies using JETSET [98] show that for this data, the extracted dN/dy_T is about 10% larger than $dN/d\eta_T$ for $|y_T| \sim 0$ and about 10% smaller than $dN/d\eta_T$ for $|y_T| \sim 4$.

The similarity in the integrated multiplicity as well as the shapes of pseudorapidity distributions for e^+e^- and the most central Au+Au data suggests that there should be a similarity in the evolution of the mid-rapidity density with collision energy, an expectation that is verified by the data (see Fig. 19). This additional close correspondence between the properties of central Au+Au and e^+e^- multiplicity data suggests that the agreement results from some underlying feature of particle production as opposed to being an accidental coincidence. In particular, an understanding of why the shape of the pseudorapidity distribution for Au+Au collisions approaches that of e^+e^- for more central interactions might prove particularly enlightening.

The arguments presented in Sect. IV A concerning total charged particle multiplicities should not be interpreted to imply that all observables in A+A will match




FIG. 17: The main panel shows the distributions of pseudo-rapidity densities of charged particles emitted in d+Au collisions with $\sqrt{s_{NN}}=200$ GeV at various centralities (see Fig. 15) divided by the distribution for inelastic $\bar{p}+p$ collisions at the same energy [85]. The positive pseudorapidity direction is that of the deuteron. Centralities are labeled by the fraction of total inelastic cross section in each bin, with smaller numbers being more central. The lower and upper arrows on the right show the average number of participants in the deuteron for the most peripheral (80–100%) and most central (0–20%) bin, respectively. The inset shows the values averaged over several bins in negative pseudorapidity plotted versus the average number of participants in the Au nucleus for each centrality bin.

those in p+p at a factor of two higher \sqrt{s} . The mid-rapidity particle densities provide an instructive counterexample. Since the same total number of particles in p+p at a higher \sqrt{s} are distributed over a broader range of pseudorapidity (see, for example, the left panel of Fig. 21), a factor of two shift in the p+p center-of-mass energy obviously cannot result in mid-rapidity densities equal to those measured in A+A. An examination of Fig. 19 reveals that the data confirm this expectation.

A less trivial counterexample is illustrated in Fig. 20 which shows ratios of the yields of antiprotons over protons emitted near mid-rapidity in p(\bar{p})+p and A+A collisions as a function of $\sqrt{s_{NN}}$. The ratios for d+Au at $\sqrt{s_{NN}}=200$ GeV (discussed in Sect. II B 1 and shown in Fig. 5) are consistent with the value shown on the figure for p+p. As discussed in Section II B, the relevant physics for understanding this ratio involves the interplay of baryon transport and antibaryon-baryon pair creation. In this case, in contrast to the situation for particle multiplicities, it is clear that the ratios for the nucleus-nucleus data are comparable to those in nucleon-nucleon colli-

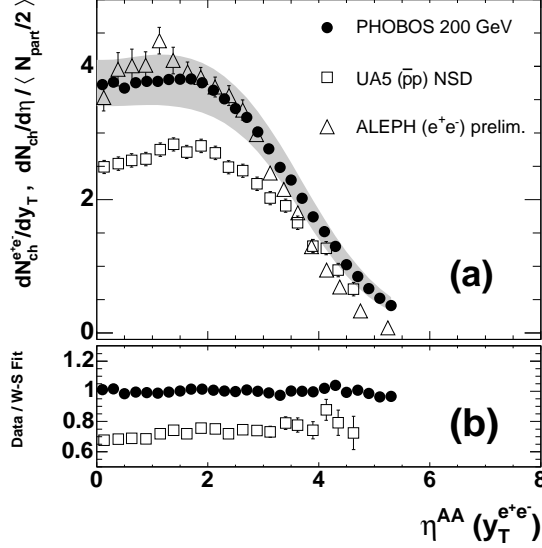


FIG. 18: The dN/dy_T distribution for charged particles emitted in e^+e^- collisions is compared to the $dN/d\eta$ distribution for charged particles emitted p+p and the normalized $dN/d\eta$ distribution for charged particles emitted in the 3% most central Au+Au collisions. All three systems are at $\sqrt{s_{NN}}$ or \sqrt{s} of 200 GeV. In the bottom panel, the Au+Au and p+p data are both shown divided by a fit to the former.




FIG. 19: Pseudorapidity particle density near mid-rapidity as a function of energy for p+p, A+A and e^+e^- reactions (where the e^+e^- density is dN/dy_T , as explained in the text). Nucleus-nucleus data, shown for central collisions, have been divided by the number of participating nucleon pairs. Note that mid-rapidity particle densities are not available for lower energy p+p or e^+e^- collisions, in the latter case due to the lack of a well defined jet structure.

sions at significantly *lower* center-of-mass energies. Although this result may not be unexpected given the larger baryon rapidity loss in A+A as compared to p+p, it serves to illustrate the importance of a systematic study to unravel the dynamical differences between the simpler and more complicated systems.

Finally, the extraction of nuclear modification factors,




FIG. 20: Antiproton to proton ratios near mid-rapidity as a function of $\sqrt{s_{NN}}$ for p+p collisions (open symbols) and central A+A collisions (filled symbols). Error bars include both statistical and systematic errors.

R_{AA} , requires the explicit use of a p+p reference spectrum. In light of the discussion presented above, one might consider comparing Au+Au transverse momentum spectra to those extracted from p+p collisions at a higher collision energy. Alternatively, p+p data at the same energy, but for events selected to have an average multiplicity matching the multiplicity per participant pair in A+A, could be used. At \sqrt{s} values of 200 GeV and above, it is known that either of these choices results in a p_T spectrum that falls off less steeply [99–101]. One should keep these ambiguities in mind when interpreting data for R_{AA} or $R_{AA}^{N_{part}}$. These issues do not affect modification factors such as $R_{PC}^{N_{part}}$ which compare A+A at different centralities.

Of course, for very peripheral A+A collisions, all observables must evolve to match those in p+p (or, to be exact, the appropriate mix of p+p, p+n, and n+n) collisions at the same \sqrt{s} . The current PHOBOS analysis of Au+Au collisions typically spans a range of impact parameters corresponding to a variation in the average number of participants in centrality bins of more than a factor of 5–6, i.e. from roughly 60 up to 350 or more. One remarkable aspect of this broad data set is that, over this range, the total particle multiplicity deviates very little from its central value when suitably normalized by the number of participants (see Fig. 12). The normalized pseudorapidity density near mid-rapidity does vary and is tending towards the p+p value but is still far above it for the most peripheral collisions studied to date (see discussion in Sect. IV F). The shape and magnitude of the transverse momentum distributions also vary but only slightly and they show little sign of tending towards the p+p distribution (see Fig. 8). One can speculate that these deviations between peripheral Au+Au and p+p collisions might result from the fact that the number of collisions per participant (or the fraction of the participants that are multiply struck) rises extremely rapidly with decreasing impact parameter for these most grazing collisions (See Appendix B 1 and Fig. 36).

In summary, there may not be a single data set, such

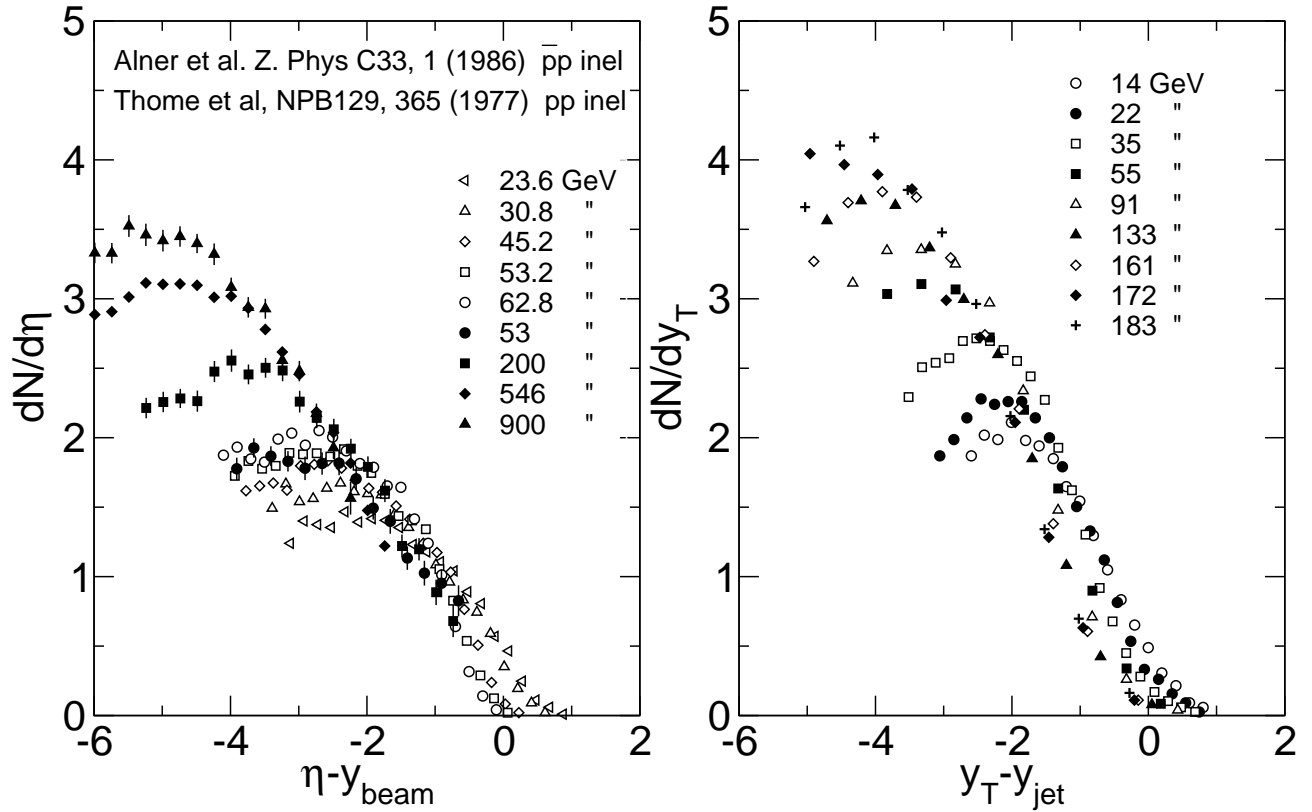


FIG. 21: (Left Panel) Distributions of pseudorapidity density of charged particles emitted in $p(\bar{p})+p$ collisions at a range of energies versus the variable $\eta - y_{beam}$. (Right Panel) Similar data for particles emitted along the jet axis in an e^+e^- collision versus the variable $y_T - y_{jet}$, defined in Appendix B 2. In both cases, when effectively viewed in the “target” rest frame, these collisions exhibit longitudinal scaling (energy independence).

as $p+p$ data at the same or a different $\sqrt{s_{NN}}$, which can serve as an ideal “reference” set for characterizing the complete dynamics of $A+A$ interactions. Instead, the properties of a broad range of systems should be studied over a range of energies and centralities to elucidate the similarities and differences among them. This should lead to a more complete understanding of the salient features of the underlying physics.

E. Extensive Longitudinal Scaling

This section describes several features of the pseudorapidity dependence of observables in a variety of systems. In particular, the distributions of particle yield and elliptic flow are found to be largely independent of center-of-mass energy over a broad region of pseudorapidity when viewed in the rest frame of one of the colliding particles. In addition, no evidence is found for a broad region near mid-rapidity displaying the characteristics expected for a boost-invariant scenario.

1. Longitudinal dependence of particle production: Elementary Systems

Before considering the energy dependence of pseudorapidity distributions in heavy ion collisions, it is instructive to review the extensive history of literature devoted to interpretations of, and expectations for, such distributions in simpler systems. A very general picture of elementary hadron-hadron collisions emerged in the late 1960’s, consisting of two sources of particle production. This concept led to the prediction of two types of scaling laws for the distributions of final state particles in different regions of the longitudinal momentum space.

Particles near beam and target rapidity were thought to be governed by the “limiting fragmentation hypothesis” [102]. In this model, the momentum distribution of particles of species “i” in the rest frame of one of the original colliding hadrons (commonly denoted with a prime to distinguish it from the center-of-mass frame), $E_i d^3N_i/dp'^3$, or equivalently $d^3N_i/dy' dp_T d\phi$, becomes energy-independent at high enough collision energy. The central concept is that the “projectile” hadron, when seen in the frame of the “target”, is Lorentz-contracted into a very narrow strongly-interacting pancake which passes through the target. This interaction leaves behind a com-

plicated excited state whose properties do not depend in detail on the energy or even identity of the projectile, and which then “fragments” into a final state distribution of particles, $E_i d^3 N_i / dp^3$. It was generally assumed that this process produced particles primarily in a restricted window of rapidity around $y'=0$, possibly even leading to a complete lack of particles at mid-rapidity in a very high energy hadron-hadron collision [103].

Particles near mid-rapidity in the center-of-mass frame, by contrast, were expected to form a rapidity plateau with a constant dN/dy , independent of energy and the nature of the hadrons in the initial collision [104, 105]. Similarly, in heavy ion collisions, a boost-invariant central plateau where “the initial conditions ... are invariant with respect to [longitudinal] Lorentz transformations” (i.e. observables are independent of y) was predicted [46]. Furthermore, the extent of this boost-invariant region was expected to grow with energy.

For elementary collisions such as p+p, and even e^+e^- , this general picture failed completely. Instead, the extensive longitudinal scaling, seen in the form of x_F scaling, pointed the way to the current view in terms of QCD, modeled for instance in the widely used Pythia code [106]. This formulation generalized the concept of “fragmentation”, which “describes the way the creation of new quark-antiquark pairs can break up a high-mass system into lower-mass ones, ultimately hadrons” [107]. It should be noted that energy independence, or scaling, in $E_i d^3 N / dp^3$ (i.e. full “limiting fragmentation”) implies scaling of both dN/dy' and dN/dx_F .

Figure 21 shows $dN/d\eta'$ for $p(\bar{p})+p$ collisions [108, 109] and $dN/d(y_T - y_{jet})$ for e^+e^- collisions [110] (see Appendix B 2 for definitions). For these elementary systems, instead of a growing boost-invariant plateau, an extended version of limiting fragmentation is found, which leads to longitudinal scaling (energy independence) over more than four units of rapidity, extending nearly to mid-rapidity. The entire system can be described in terms of either string “fragmentation” or in terms of a parton cascade, leading naturally to extensive longitudinal scaling.

2. Longitudinal dependence of particle production: d+A and p+A

In the case of asymmetric systems, the concept of limiting fragmentation can be explored separately in the rest frame of the two projectiles. Such studies, applied to hadron-nucleus collisions, were of particular interest in the 1970's [91]. The specific question was whether the region of rapidity in which the particle yield is A-dependent expands with increasing collision energy [111–113]. Many models predicted that an extensive A-dependent region, indicative of long-range order, should not occur. Instead, only a localized region near the rapidity of the larger collision partner would be affected by the target mass, and further, the height and width of this region was expected to be independent of, or at most weakly




FIG. 22: A compilation of distributions of pseudorapidity densities of charged particles emitted in p+A and d+A collisions at a variety of energies. Grey tracks are included in the distributions shown for emulsion data. As in Fig. 24, the data are plotted versus the variables $\eta + y_{target}$ and $\eta - y_{beam}$ calculated using the rapidity of the larger (left panels) or smaller (right panels) of the colliding species. Note that, as in Fig. 24, the data at all energies and at both ends of the pseudorapidity range follow common curves.

dependent on, beam energy. One prediction of these expectations was that the integrated yield in p+A would approach the value observed in p+p at high beam energies, since the small A-dependent region would become increasingly unimportant [114]. Instead, to the surprise of many people, a broad A-dependent region was observed, displaying characteristics very similar to the extensive longitudinal scaling observed in simpler systems [84, 86, 93, 95, 115, 116].

Pseudorapidity distributions from PHOBOS can be used to extend these studies to d+A collisions at RHIC energies. In Fig. 22, a compilation of pseudorapidity density data for proton+(nuclear emulsion) [115, 116] and p+Pb [86] at various energies is shown, together with PHOBOS data for d+Au at $\sqrt{s_{NN}}=200$ GeV [81], with the centrality and normalization for the d+Au results chosen appropriately. To be more specific, the d+Au pseudorapidity densities are divided by the number of participating nucleons in the deuteron (by definition this would be unity for p+A). Furthermore, the d+Au centrality bin was selected such that the ratio of the number of participating nucleons in the Au nucleus to the number in the deuteron was equal to the number of participating nucleons from the lead or emulsion in p+A. This latter quantity is commonly denoted ν , the average

number of collisions per participant in the smaller projectile (see definitions in Appendix B 1). Fig. 22 clearly demonstrates that extensive longitudinal scaling (limiting fragmentation) also is manifested in d+A collisions at RHIC energies.

3. Longitudinal dependence of particle production: Au+Au at RHIC

The uniquely broad pseudorapidity coverage of the PHOBOS detector allows similar studies to be performed for heavy ion collisions at RHIC energies. At first the pseudorapidity distributions themselves, shown in Fig. 1, suggest that $dN_{ch}/d\eta$ may develop a small boost-invariant central plateau, but these plots are misleading for this purpose. Pseudorapidity is known to distort the rapidity distribution for production angles near 0° and 90° . Demonstrating this point, the rapidity distributions of positive pions measured by BRAHMS [117], as well as similar data at lower energies [75, 118], are all well represented by Gaussian fits, as shown in Fig. 23. In short, there are no indications of the existence of a broad boost-invariant central plateau in the final particle distributions.

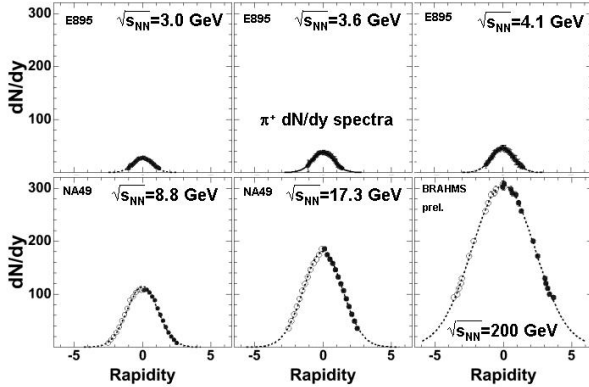


FIG. 23: Rapidity densities of positive pions emitted in central collisions of Au+Au (AGS and RHIC) [117, 118] and Pb+Pb (SPS) [75] at a variety of beam energies. Note that in contrast to Fig. 1, yields in rapidity space are well represented by Gaussians with no evidence for a broad mid-rapidity plateau.

In Fig. 24, the data shown in Fig. 1 are shifted to effectively plot them in the rest frame of one of the gold nuclei. Lorentz boosts of pseudorapidity, η , are not as trivial as those of rapidity, but $\eta' \equiv \eta - y_{beam}$ (or $\eta + y_{beam}$) approximates y' . Furthermore, as noted above, limiting fragmentation implies scaling in the full distribution, $E_i d^3 N_i / dp^3$. Since η' is just a function of (y', p_T, m_i) , limiting fragmentation also directly implies scaling in $dN/d\eta'$ as well. The data at both centralities show an extensive scaling with the longitudinal velocity




FIG. 24: Distributions of pseudorapidity densities of charged particles emitted in Au+Au collisions at three energies and two centrality ranges are plotted versus $\eta' \equiv \eta - y_{beam}$ (or $\eta + y_{beam}$). In the far right panel, data for positive and negative η have been averaged to generate data versus $|\eta| - y_{beam}$. Systematic errors are not shown. Note that the data from all three energies follow a common curve.

in the rest frame of one of the projectiles, identical behavior to that seen in simpler systems (see, for example, [85, 87, 88, 95]). Similar behavior in nucleus-nucleus collisions over a somewhat narrower range in η' was first observed by BRAHMS [119, 120].

Figure 24 illustrates one example of how the scaling behaviors can be used to infer the properties of particle production which lie outside the experimental acceptance at large collision energies. If one accepts the assumption that the η' distributions at all energies are identical in the region corresponding to larger η , the data from lower energies can be used to constrain the extrapolation of the higher energy data to the full phase space. In addition, it should be noted that the corrections to the PHOBOS multiplicity data depend strongly on emission angle of the particles and also are significantly asymmetrical between positive and negative pseudorapidities. The latter effect results primarily from the offset of the PHOBOS magnet from the center of the interaction region (see Fig. 33). The good agreement seen when comparing particles emitted at different angles and for both signs of pseudorapidity indicates the robustness of the analysis procedure as well as providing interesting physics insight.

Fig. 24 illustrates the observation that longitudinal scaling holds over an even more extensive range of pseudorapidity in these seemingly complex high energy A+A collisions at RHIC. Based on the pseudorapidity distribution (and, as will be discussed in following sections, elliptic flow and perhaps even HBT), no evidence is seen in any hadron-hadron or ion-ion collisions for two energy independent fragmentation regions separated by a boost invariant central plateau which grows in extent with increasing collision energy. Thus, the expectation from

the boost-invariant description of the energy evolution of rapidity distributions is not valid for heavy ion collisions either. In fact there is no boost invariant central plateau and, instead, the rapidity distribution appears to be dominated by two “fragmentation” regions, whose extent increases with energy. We call this effect “extensive longitudinal scaling”.




FIG. 25: Pseudorapidity dependence of elliptic flow of charged particles for the 40% most central collisions of Au+Au (average number of participating nucleons indicated) at a variety of beam energies. Note the linear fall-off at higher $|\eta|$ and the lack of evidence for a constant value over a broad mid-rapidity region. Boxes indicate systematic uncertainties (90% C.L.).

4. Longitudinal dependence of elliptic flow: Au+Au at RHIC

In addition to the pseudorapidity distributions of yields of produced particles, longitudinal scaling can also be seen in the elliptic flow of particles produced in heavy ion collisions. As discussed in Section III, the elliptic flow parameter, v_2 , provides a sensitive probe of the properties in the early stages of the collision, one of which is the presence or absence of boost-invariance. Boost invariant “initial conditions” (i.e. right after the collision) should lead to a boost-invariant $v_2(y)$. Kinematic effects result in a difference between $v_2(y)$ and $v_2(\eta)$ [121], leading to an increase of about 10–15% near mid-rapidity in the latter. However, the small magnitude of this difference means that it does not affect the conclusions discussed here and that a boost-invariant scenario (in rapidity) should also result in elliptic flow which is approximately flat over a large region of pseudorapidity. In Fig. 25, the pseudorapidity dependence of the elliptic flow parameter, v_2 , is shown for semi-central Au+Au events at various energies [122]. Clearly, no boost invariant central plateau is seen. Thus, there are no indications of the existence of a broad boost-invariant central plateau in the final particle distributions or in the state formed shortly after the collision as reflected by v_2 .

In Fig. 26, the elliptic flow data from Fig. 25 is replotted

effectively in the rest frame of one of the gold nuclei. Once again the phenomenon of “limiting fragmentation” is revealed, this time for v_2 [122]. As discussed above, there is a small modification of the shape if v_2 is plotted versus rapidity instead of η but this change does not significantly impact the comparison of different energies. There appears to be a single universal curve governing the elliptic flow as a function of η' over a broad range down to mid-rapidity at each energy studied. The limiting fragmentation behavior of elliptic flow in Fig. 26 has further implications since elliptic flow builds up early in the collision. Therefore, the dependence on the location in $|\eta'|$ space must reflect the conditions very shortly after the collision, and then these early conditions lead to the measured elliptic flow.

5. Longitudinal dependence: Lessons from HBT

Particle interferometry, in the form of Hanbury-Brown Twiss (HBT) correlations [123, 124], provides an extra, although much more indirect, test of the ideas of boost-invariance in heavy ion collisions. Since pions are bosons, they constructively interfere when they are near each other in phase space. Correlation measurements in momentum space can therefore reveal the source size in position space. In particular, HBT correlations are sensitive to the spatio-temporal distributions of particles at thermal freeze-out (i.e. the point of the last elastic interactions). Appendix B 3 contains more details including a description of the source parameterizations. Most theoretical studies of HBT assume ideal (i.e. non-viscous) hydrodynamics and a boost-invariant source which exhibits longitudinal Hubble flow ($z = v_z t$, where z and v_z are the longitudinal position and velocity, respectively). These assumptions simplify the coupled differential equations and allow the use of 2D transverse expansion overlaid on the boost-invariant longitudinal expansion (a scenario often called 2+1D hydrodynamics). While this basic hydrodynamic picture was roughly successful in describing some aspects of the elliptic flow (see Figs. 6 and 7), these models have failed to describe the HBT data from RHIC [125–127].

Figure 27 shows the results of fits using the Bertsch-Pratt parameterization along with the R_o/R_s ratio from $\sqrt{s_{NN}}=62.4$ and 200 GeV Au+Au collisions [130]. Under the assumptions of boost-invariant hydrodynamics, the R_o/R_s ratio should be large if a long-lived source is formed and should typically be larger than $\sqrt{2}$ in any case. By contrast, the ratio of R_o/R_s appears to be closer to unity in heavy ion collisions. Similar results were found in heavy ion collisions at lower energies (see references in [130]). The smallness of both R_o/R_s and R_l has come to be known as the “HBT puzzle”. It has been postulated that relaxing the assumption of boost-invariance [131, 132] or allowing non-zero viscosity [133] may resolve this discrepancy.

The detailed nature of the longitudinal properties of




FIG. 26: The flow data of Fig. 25 are shown in the left (center) panel versus the variable $\eta' = \eta - y_{beam}$ ($\eta' = \eta + y_{beam}$). In the right panel, data at positive and negative pseudorapidity have been averaged to give v_2 as a function of $|\eta|$. These results were then plotted versus the variable $\eta' = |\eta| - y_{beam}$. As for the particle densities shown in Fig. 24, the flow data at all energies follow a common curve. In the case of flow, this curve holds over the entire range from beam or target to mid-rapidity.

particle production can also be explored by HBT measurements, in this case in a very direct way as shown in Fig. 28. The data show the average rapidity of the source of the pions (derived from the source velocity in the Yano-Koonin-Podgoretskii parameterization) as a function of the rapidity of the pions themselves [130]. A clear systematic trend is observed, and again the results are very similar to what was found at the SPS [134]. Under the simple assumption of all pions being emitted from a single source located at the center of mass, the ordinate of all points would be equal to zero. If, instead, the system consisted of a series of independent sources at different rapidities (i.e. a strong longitudinal position-momentum correlation) the points would fall on the line. The “locality” revealed by HBT studies of pion correlations in rapidity space suggests that the longitudinal distribution of particle properties is established very early, with the subsequent evolution and freezeout having only short range correlations in rapidity.

6. Extensive Longitudinal scaling: Summary

To summarize this section, the data demonstrate that extensive longitudinal scaling, reminiscent of “limiting fragmentation” over an extended region of longitudinal phase space, seems to be a dominant feature of particle production for all colliding systems. Based on all of the data, no evidence is seen in any hadron-hadron or ion-ion collisions for two energy independent fragmentation regions separated by a boost invariant central plateau which grows in extent with increasing collision energy. The lack of a broad boost-invariant central plateau is seen in both the final particle distributions and in the state formed shortly after the collision as reflected by v_2 . It is difficult to reconcile this with the common assumption that particle production at mid-rapidity results from different physics than that in the fragmentation region, particularly at the higher energies. Furthermore,

the similarity of the longitudinal scaling of both particle densities and elliptic flow suggests the possibility of some direct connection between the two, implying that the final particle multiplicities also result from the properties of the very early evolution.

A good way to appreciate the significance of these results is to consider what would be observed in the detectors if a collider could operate its two beams at different energies. For simplicity, the conventional RHIC designation for the two counter-rotating beams, namely “blue” and “yellow”, will be used. If the energy of the blue beam was set to a rapidity of 2, for example, the results show that, as the rapidity of the yellow beam was increased up to a little beyond 2, the particle density and elliptic flow seen in the detectors covering the blue beam hemisphere would show a gradual increase and then reach a limiting value. With the blue beam fixed at a rapidity of 2, the particle density would not increase beyond this limiting value on the blue beam side even if yellow was set to infinite rapidity. The only way to further increase the particle density or elliptic flow in the blue beam hemisphere would be to increase the energy of the blue beam.

In p+p collisions, extensive longitudinal scaling was understood to be a consequence of x_F scaling in string fragmentation (or, equivalently, in parton cascades). No similar, widely accepted, explanation exists for limiting fragmentation in the more complex p+A, d+A, and A+A collisions.

F. Factorization of Energy and Centrality dependence

The previous sections have described separately the dependencies of a variety of observables on energy and centrality. These independent discussions may have obscured the remarkable extent to which these two dependencies factorize. This section will describe several aspects of PHOBOS data which display this phenomenon.




FIG. 27: Bertsch-Pratt parameters R_o , R_s , R_l and R_o/R_s for $\pi^+\pi^-$ pairs emitted in central collisions of Au+Au at $\sqrt{s_{NN}}$ of 200 GeV (left panels) and 62.4 GeV (right panels) as a function of pair transverse momentum k_T . For comparison, data from STAR [128] (open stars) and PHENIX [129] (open crosses) are presented at $\sqrt{s_{NN}}=200$ GeV. PHOBOS systematic errors are shown as boxes; systematic errors from STAR and PHENIX are not shown.

One simple example of factorization was revealed by the PHOBOS measurements of the total charged particle multiplicity divided by the number of pairs of participating nucleons in Au+Au collisions at three energies, from 19.6 to 200 GeV (see Fig. 12). One sees that the data for the different energies are separated by a factor that is constant as a function of centrality. In other words, the centrality and energy dependence of the yield per participant in Au+Au collisions factorize over the range of the two control variables. In this case, the factorization occurs trivially, as the total charged particle yield per participant is centrality-independent at all energies. Whether this factorization is a fundamental property of particle production in Au+Au collisions can be tested by studying the yields per participant more differentially in pseudorapidity and transverse momentum.

In Fig. 29, the pseudorapidity density of charged particles per participant pair near mid-rapidity is shown as a function of centrality for collision energies of 19.6 GeV and 200 GeV [26]. Over the centrality range shown here, the normalized yield at mid-rapidity increases by approx-




FIG. 28: The rapidity source parameters for $\pi^+\pi^+$ pairs emitted in central Au+Au collisions at RHIC and Pb+Pb collisions at the SPS. This element of the Yano-Koonin-Podgoretskii parameterization specifies the rapidity (in the nucleus-nucleus center-of-mass system) of the source, Y_{YKP} , from which the pions were emitted. The abscissa of this plot is the average rapidity of the pions themselves. The filled symbols are the measured data, the open symbols have been reflected about mid-rapidity. The line with a slope of 1 is drawn to guide the eye.

imately 25% from mid-peripheral to central collisions. Early theoretical explanations attributed this increase to the contribution of the hard component of particle production, which would grow with the relative increase in the number of binary nucleon-nucleon collisions in more central events. As an example of such a superposition of soft and hard particle production, the results of a HIJING calculation are shown as dashed lines. The model shows an increase in the yield per participant pair, although steeper than that seen in the higher energy data.

However, this explanation is challenged by the detailed study of the energy dependence of mid-rapidity particle yields shown in Fig. 30, where the centrality dependence of the ratio of 200 GeV to 19.6 GeV data is plotted. Within the experimental uncertainty, this ratio is independent of centrality, whereas the contribution from hard processes would be expected to show a large increase over this collision energy range. This is illustrated by the HIJING prediction for this ratio (shown as a dashed line), which completely fails to capture the factorization of energy and centrality dependence for the mid-rapidity yield per participant. A similar result was found earlier (over a smaller span in beam energy) using the centrality dependence of normalized mid-rapidity yields from Au+Au at $\sqrt{s_{NN}}=130$ GeV [135, 136].

The results of an attempt to investigate the interplay




FIG. 29: Pseudorapidity density of charged particles emitted near mid-rapidity divided by the number of participant pairs as a function of number of participants. Data are shown for Au+Au at collision energies of 19.6 and 200 GeV. The grey ellipses show the 90% C.L. systematic errors. The results of two models and one parameterized fit are shown for comparison.

of hard and soft scattering without invoking a complicated model are shown as dotted lines. In this case, a very simplistic two component fit [137] was performed to separately extract the fractions of the particle yield which scaled with the number of participants (soft scattering) and the number of collisions (hard scattering). A reasonably good fit to the data is found but the fitted parameters suggest that, within the uncertainties, there would be an identical contribution from hard scattering at both beam energies, a result which is totally unexpected for minijet dominated physics.

Also shown in Fig. 30 is the result of a saturation model calculation [32, 138]. This model, which, as mentioned in Section II A, yields a reasonably good match to the energy evolution of particle yields at RHIC energies, also does a much better job of describing the centrality evolution than the HIJING model.

Another example of non-trivial centrality dependence that is energy independent was shown by the pseudorapidity distributions in Figs. 14 and 24. The former showed that the shape of the distributions differed significantly as a function of centrality. The latter demonstrated that the distributions at different beam energies were found to line up when plotted in the approximate rest frame of one of the incoming nuclei, i.e. using the variable $\eta' \equiv \eta - y_{beam}$. Thus, the shape evolution with centrality is independent of beam energy over a very broad range in η' .

Additional evidence for factorization is provided by the transverse momentum distributions briefly mentioned in Sect. III. In the absence of medium effects, one would expect that the volume scaling (i.e. proportionality to N_{part}) observed for the bulk production of hadrons turns




FIG. 30: Ratio of the pseudorapidity densities of charged particles emitted near mid-rapidity for Au+Au at 200 GeV over 19.6 GeV as a function of the number of participants. The error bars include both statistical and 1- σ systematic errors. The ratios for the same two models and one fit shown in Fig. 29 are displayed for reference.

into scaling with the number of binary collisions (N_{coll}) when measuring reaction products of point-like hard processes. This transition should be visible when studying particle production as a function of transverse momentum. However, as is now known (see Fig. 8), particle production at large transverse momenta seems to be strongly modified in the presence of the medium in heavy-ion collisions. The strength of this modification is more clearly illustrated in Fig. 31 which shows the nuclear modification factor for charged hadrons in six bins of p_T as a function of N_{part} . In the figure, yields at a given transverse momentum in collisions of varying centrality were normalized by the number of participant pairs and then divided by a fit to the same quantity in central data (see Appendix B 3 for definition). Data for p+p collisions from UA1 [100] are shown with the same normalization factor. It is striking to see that the medium modification results in charged particle yields that, over the centrality range studied here, more closely scale with N_{part} than with the number of binary collisions, even for transverse momenta above 4 GeV/c.

The observation of N_{part} scaling at high transverse momentum suggests that the medium is almost completely “black” or “absorbing” to produced fast particles. This conclusion follows if one assumes N_{coll} scaling of the primary production throughout the entire volume of the collision zone followed by complete absorption except on the surface. The volume to surface ratio (proportional to the nuclear radius R or equivalently $A^{1/3}$) has a similar centrality dependence to the ratio of the number of collisions and participants. However, since the centrality dependence of particle production is seen to be very similar at all transverse momenta, it is also possible that the usual

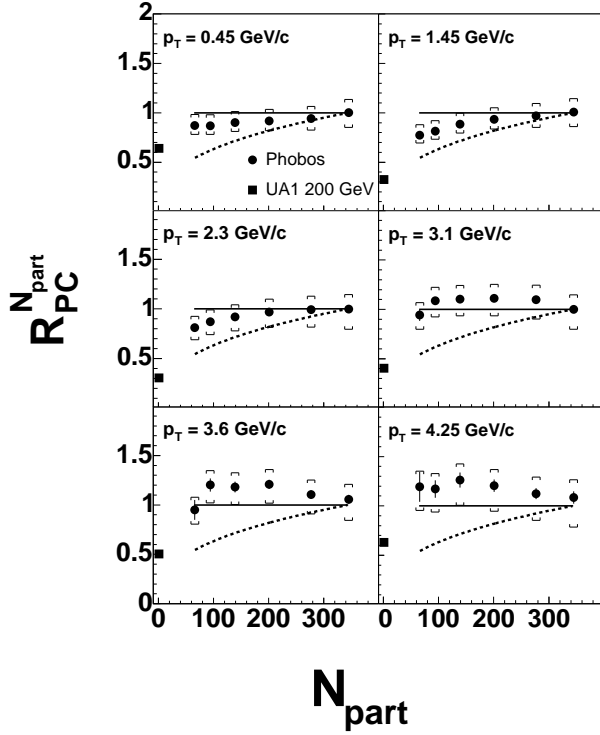


FIG. 31: Particle yield normalized by the number of participant pairs and then divided by central data (identical to $R_{PC}^{N_{part}}$ defined in Appendix B 3) as a function of centrality for Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV, for six transverse momentum ranges. Bars and brackets show statistical and systematic uncertainties, respectively. The solid (dashed) line shows the expectation for N_{part} (N_{coll}) scaling from peripheral to central collisions. Squares show data for p+p collisions from UA1 [100] with the same normalization factor.

simplistic assumption of participant dominance at low p_T evolving into collision dominance at higher values needs to be reconsidered.

Data from the most recent RHIC run have been used to study the evolution of the transverse momentum distributions as a function of both collision centrality and energy. The measurements were performed near mid-rapidity at collision energies of 62.4 GeV and 200 GeV [64]. In Fig. 32, particle production as a function of centrality and p_T is shown for these two energies in terms of $R_{AA}^{N_{part}}$ and $R_{PC}^{N_{part}}$. $R_{AA}^{N_{part}}$ shows the variation in the yield per participant pair relative to p+p collisions (upper row of Fig. 32) and $R_{PC}^{N_{part}}$ shows the variation in yield per participant pair relative to central Au+Au collisions (bottom row).

As discussed earlier, the range in p_T from a few hundred MeV/c to more than 4 GeV/c is believed to cover very different regimes of particle production, from soft coherent processes to independent binary scattering. Over the collision energy range from 62.4 to 200 GeV, overall particle production in p+p increases by less than a factor

2, whereas the yield at $p_T=4$ GeV increases by an order of magnitude. This clearly shows the change in the balance of lower and higher transverse momenta particles, believed to reflect the different energy dependencies of soft and hard particle production in p+p collisions over this energy range. For central Au+Au collisions however, the ratio of the yields between 200 GeV and 62.4 GeV at $p_T=4$ GeV/c in central Au+Au is only about 4 (with a factor of 1.6 increase in the p_T -integrated multiplicity), i.e. the huge increase in the yield of high p_T particles in p+p is not reflected in Au+Au.

The top row of Fig. 32 clearly demonstrates that the shape and magnitude of the transverse momentum spectra depend strongly on beam energy and to a lesser extent also on centrality. In particular, at both energies the yield per participant at any given p_T changes by less than 25% over the centrality range from 60 to 340 participants, with an even smaller variation at the highest p_T . Even more surprisingly, the comparison in terms of $R_{PC}^{N_{part}}$ in the bottom row of the figure shows that the remaining variation of the yield per participant pair is the same for both energies over the full p_T and centrality range. This means that the energy and centrality dependence of particle production also factorizes in p_T over this energy range. This is particularly striking, as the factorization therefore covers both the bulk particle production at low p_T , as well as rare particle production at intermediate and high p_T , believed to be governed by different particle production mechanisms. In particular at intermediate p_T above 1 GeV, particle production is thought to be influenced by the effects of radial hydrodynamic flow, the p_T broadening due to initial and final state multiple scattering (“Cronin effect”), the balance between ‘soft’ and ‘hard’ particle production, recombination and fragmentation, and the in-medium energy loss of fast partons. All of these contributions to the overall particle yields are expected to show distinctly different centrality and energy dependencies at different p_T , yet the overall result is a factorization of energy and centrality dependence at all p_T within the experimental uncertainty. This result, combined with similar observations for the rapidity distributions, strongly suggests that the data reflect the dominant influence of yet-to-be-explained overall global constraints in the particle production mechanism in A+A collisions.

V. CONCLUSION

PHOBOS data and results from the other RHIC experiments, combined with very general arguments which are either model independent or depend on fairly simple model assumptions, lead to a number of significant conclusions.

In central Au+Au collisions at RHIC energies, a very high energy density medium is formed. Conservative estimates of the energy density at the time of first thermalization yield a number in excess of 3 GeV/fm³, and the




FIG. 32: Nuclear modification factors versus transverse momentum for Au+Au at two beam energies and a variety of centralities calculated using two different reference distributions; $N_{part}/2$ (top row) times p+p yields, or the ratio of N_{part} (bottom row) times the distribution for central Au+Au. Filled symbols are for $\sqrt{s_{NN}}=200$ GeV, open symbols are for 62.4 GeV. Centralities are labeled by the fraction of total inelastic cross section in each bin, with smaller numbers being more central and the number of participants at the lower energy are indicated. Note the small variations in both the magnitude and shape of the ratios calculated using N_{part} and also that R -factors normalized using central Au+Au data (bottom row) are identical at the two beam energies.

actual density could be significantly larger. This is far greater than hadronic densities and so it is inappropriate to describe such a medium in terms of simple hadronic degrees of freedom. Unlike the weakly interacting QGP, naively imagined by a large part of the community before RHIC turn-on, the produced medium was found to be strongly interacting. If this medium is a new form of QCD matter, as one would expect from lattice gauge calculations for such a high energy density system, the transition to the new state does not appear to produce any signs of discontinuities in any of the observables that have been studied. To the precision of the measurements, all quantities evolve smoothly with energy, centrality, and rapidity. This would suggest that, if there is a phase change at RHIC, the transition from this strongly interacting high energy density medium to a hadronic gas is a cross-over.

An equally interesting result was the discovery that much of the data can be expressed in terms of simple scaling behaviors. In particular, the data clearly demonstrate that the number of participating nucleons, N_{part} , is a key parameter which determines much of the phenomenology. Further, the total particle yields per participant from different systems are close to identical when compared at the same available energy; the longitudinal velocity dependences of elliptic flow and particle yield are energy independent over a very broad range, when viewed in the rest frame of one of the colliding nuclei; and many characteristics of the produced particles factorize to a surprising degree into separate dependences on centrality and beam energy.

All of these observations point to the importance of the geometry of the initial state and the very early evolution of the colliding system in determining many of the properties of the final observables. Future data at RHIC, most especially collisions of lighter nuclei, as well as higher energy nucleus-nucleus data from the LHC, will help to further evaluate the range of validity of these scaling behaviors. It is possible that models which describe the initial state in terms of a CGC will play a role in explaining some or all of these scaling properties, but such an identification is not yet clear. What is clear is that these simple scaling features will constitute an integral component or essential test of models which attempt to describe the heavy ion collision data at ultrarelativistic energies. The unifying features may, in fact, provide some of the most significant inputs to aid the understanding of QCD matter in the region of the phase diagram where a strongly interacting, very high energy density medium is created.

Acknowledgments

The PHOBOS collaboration would like to express our gratitude to the RHIC installation and management personnel and to the operations staff of the BNL Collider-Accelerator Department for their hard work and truly spectacular success. Without the broad range of colliding species and energies, and continually increasing luminosities, provided by RHIC over a very short span of time, the PHOBOS physics program would not have

been nearly so rich. We would like to express special thanks to S.Basilev, B.D.Bates, R.Baum, A.Białas, M.Ceglia, Y.-H.Chang, A.E.Chen, T.Coghen, C.Conner, J.Corbo, W.Czyż, B.Łabrowski, M.Despet, P.Fita, J.Fitch, M.Friedl, K.Gałuszka, R.Ganz, J.Godlewski, C.Gomes E.Griesmayer, J.Halik, P.Haridas, A.Hayes, D.Hicks, W.Kita, J.Kotula, H.Kraner, C.Law, M.Lemler, J.Ligocki, J.Michałowski, J.Mülmenstädt, M.Neal, A.Noell, M.Patel, M.Plesko, M.Rafelski, M.Rbeiz, D.Ross, J.Ryan, A.Sanzgiri, J.Scaduto, J.Shea, J.Sinacore, S.G.Steadman, M.Stodulski, Z.Stopa, AStraczek, M.Strek, K.Surowiecka, R.Teng, B.Wadsworth, K.Zalewski, and P.Żychowski without whose efforts the PHOBOS experiment would not have been possible. We acknowledge the generous support of the Collider-Accelerator Department (including RHIC project personnel) and Chemistry Departments at BNL. We thank Fermilab and CERN for help in silicon detector assembly. We thank the MIT School of Science and LNS for financial support. This work was partially supported by U.S. DOE grants DE-AC02-98CH10886, DE-FG02-93ER40802, DE-FC02-94ER40818, DE-FG02-94ER40865, DE-FG02-99ER41099, and W-31-109-ENG-38, by U.S. NSF grants 9603486, 0072204, and 0245011, by Polish KBN grant 1-P03B-062-27(2004-2007), and by NSC of Taiwan Contract NSC 89-2112-M-008-024.

APPENDIX A: THE PHOBOS DETECTOR

The PHOBOS experimental setup is composed of three major sub-systems; a charged particle multiplicity detector covering almost the entire solid angle, a two arm magnetic spectrometer with particle identification capability, and a suite of detectors used for triggering and centrality determination. More details can be found in [139]. The active elements of the multiplicity detector and tracking detectors in the spectrometer are constructed entirely of highly segmented Si wafers with individual readout of the energy deposited in each pad [140–142]. The layout of the experiment during the 2004 run is shown in Fig. 33. An enlarged view of the region around the beam collision point is shown in Fig. 34. Table I lists the colliding systems, center-of-mass energies, and data samples collected by PHOBOS during all of the RHIC runs to date.

The Si pad detectors used to measure multiplicity consist of a single layer covering almost the entire 4π solid angle. These detectors measure the total number of charged particles emitted in the collisions as well as detailed information about their distribution in azimuthal and polar angle (or equivalently pseudorapidity, η). The Si modules are mounted onto a centrally located octagonal frame (Octagon) covering $|\eta| \leq 3.2$ as well as three annular frames (Rings) on either side of the collision vertex, extending the coverage out to $|\eta| \leq 5.4$.

The Si modules forming the arms of the spectrometer are mounted on eight frames. Depending on the trajectory, charged particles traverse between 13 and 16 layers

of Si as they pass through the spectrometer. The first layer is only 10 cm from the nominal interaction vertex. The magnet pole tips are arranged to produce almost no magnetic field in the vicinity of the first six layers. The field then rises rapidly to a roughly constant value of 2 Tesla for the remaining layers. The Si wafers are finely segmented to provide 3-dimensional space points used in the track finding. The phase space covered depends on the vertex location along the beam direction and extends over about $3/4$ of a unit of η for any given vertex location, with a total coverage of roughly $0 < \eta < 2$. Each arm covers approximately 0.1 radians in azimuth for particles that traverse all of the layers. The momentum resolution is close to 1% near 0.5 GeV/c and rises about 1% for each additional 3 GeV/c.

Particle identification is provided using two techniques. Charged particle energy loss is measured in each Si layer. Combining this information with the momentum from the tracking can separate pions from kaons out to about 700 MeV/c and from protons out to about 1.2 GeV/c. Additional particle identification is provided by two Time-of-Flight (TOF) walls each consisting of 120 plastic scintillator slats. Before the start of the 2003 RHIC run, these walls were moved farther from the interaction point, extending particle identification capability out to momenta roughly twice that achievable using energy loss in the Si. In their new locations, the TOF walls cover roughly half the azimuthal acceptance of the spectrometer.

Before the 2004 run, a small hadronic calorimeter (SpecCal) was installed behind one of the spectrometer arms. Consisting of 50 lead/scintillator modules, each 10 cm square by about 120 cm long, this detector can be used to measure the energy of high momentum particles traversing part of the spectrometer acceptance.

The primary event trigger for all colliding systems was provided by two sets of 16 plastic scintillator slats (Paddles) covering $3.2 < |\eta| < 4.5$. Imposing an upper limit on the time difference between the signal in the two arrays eliminated most beam-gas interactions and provided a rough selection of collision vertex locations along the beam line. To enhance the data sample of useful events, a more precise measure of vertex location was generated using two arrays of 10 Cherenkov counters (T0s). This was necessary because the range of vertex positions for which the multiplicity and tracking detectors have reasonable acceptance is considerably shorter than that provided by the accelerator. For different colliding systems, the T0 detectors could be moved to different locations along the beam line in order to optimize the efficiency of the vertex determination while minimizing multiple particles traversing a single counter. A more precise vertex location is found off-line using signals from the Vertex detector, which is composed of two sets of two layers each of Si modules. With high segmentation along the beam direction, correlating hits in the inner and outer layer can be used to determine the vertex along the beam line to an accuracy of better than 0.4 mm.




FIG. 33: The layout of the PHOBOS detector during the RHIC run in early 2004. The beams collide at a point just to the right of the double-dipole magnet, the top of which is not shown. The PCAL calorimeters and Zero Degree Calorimeters are drawn to scale but are located about 3 times farther from the interaction point than shown.




FIG. 34: The elements of the PHOBOS detector in the vicinity of the beam collision point.

Colliding systems, such as p+p or d+Au which produce smaller numbers of particles, have fewer events with tracks traversing the spectrometer. The spectrometer

trigger uses an additional array of scintillator slats (SpecTrig) mounted between the tracking detectors and the TOF walls. Coincidences between the SpecTrig and TOF hit slats, combined with the vertex location from T0, were used online to select events containing a high momentum track in the acceptance of both the spectrometer and the TOF.

The Zero-Degree-Calorimeters (ZDC) have a cross-sectional area of $10 \times 12 \text{ cm}^2$ centered on the direction of the beam and are located about 18 m from the nominal interaction point. Particles hitting these detectors must first traverse the initial RHIC accelerator magnet which separates the two counter-circulating beams. Therefore, the ZDC signal results almost exclusively from spectator neutrons which emerge alone (i.e. not bound in nuclear fragments) but whose transverse momentum remained close to zero after the interaction. Due to the response time of this detector, partly resulting from its long distance from the collision point, it was not possible to use ZDC signals in the primary event trigger for the bulk of the physics data. However, this device was used online in special runs to check triggering efficiency for the other detectors and also off-line in studies of centrality

RHIC Run	Colliding System and $\sqrt{s_{NN}}$	Dates of PHOBOS Data Taking	Total Events (M)
1	Au+Au 56 GeV	6/13/00–6/16/00	1.8*
	Au+Au 130 GeV	8/15/00–9/4/00	4.3
2	Au+Au 130 GeV	7/8/01	0.044
	Au+Au 200 GeV	7/20/01–11/24/01	34
	Au+Au 19.6 GeV	11/25/01–11/26/01	0.76*
	p+p 200 GeV	12/28/01–1/25/02	23
3	d+Au 200 GeV	1/6/03–3/23/03	146
	p+p 200 GeV	4/13/03–5/24/03	50
4	Au+Au 200 GeV	1/5/04–3/24/04	215
	Au+Au 62.4 GeV	3/24/04–4/2/04	22
	p+p 200 GeV	4/18/04–5/14/01	28

TABLE I: Summary of RHIC runs and data collected by PHOBOS. Note that event totals given in the last column represent the number summed over the entire variety of triggering conditions including minimum-bias, events requiring a restricted range of the collision vertex, collisions selected to be more central or more peripheral, and collisions satisfying the high- p_T spectrometer trigger. Note that triggers for the Au+Au runs at 19.6 and 56 GeV (marked with *) had very loose requirements on timing with the result that only a relatively small fraction of the events were usable in the analysis.

determination.

Similarly to the ZDC, the Proton Calorimeters (PCAL) are located behind the first accelerator magnets, but in this case next to the outer edge of one of the beam pipes. The magnets bend spectator protons to an angle of more than twice that of the beam particles so these protons will exit the beam pipe and shower in the PCAL. As with the ZDC, only individual protons, as opposed to those bound in clusters, can be detected. The PCAL is particularly useful for studies of d+Au collisions. On the side of the outgoing deuteron, the combination of the PCAL and the ZDC signals can be used to divide the event sample into p+Au, n+Au, and d+Au subsets, i.e. events in which only one or both of the incoming nucleons interacted. On the side of the outgoing Au nucleus, the PCAL is primarily sensitive to protons knocked out of the Au which is a measure of the total number of collisions suffered by the interacting nucleons in the deuteron.

APPENDIX B: DEFINITIONS OF TERMS

In this section, detailed definitions are given for the important event characterization parameters, as well as a number of the critical observables used in the physics analysis.

1. Event Characterization

In interpreting data from heavy ion collisions, the primary event characterization parameters are the energy of the collision and the overlap of the two nuclei at the moment when they interact, commonly referred to as centrality. In order to compare fixed target, colliding beam, symmetric, and asymmetric systems all on

a common footing, the collision energy is defined using the center-of-mass energy available when a single nucleon from one projectile collides with a single nucleon from the other projectile, ignoring any Fermi motion. The standard notation for this quantity is $\sqrt{s_{NN}}$, referred to as the nucleon-nucleon center-of-mass energy. For colliding beams, each of which has the same energy per nucleon, $\sqrt{s_{NN}}$ is simply twice that value. At RHIC, data have been taken for a wide range of $\sqrt{s_{NN}}$ (see Table I) ranging from a value close to the maximum achieved at the SPS up to a value more than 10 times larger.

A direct measure of the collision geometry is given by the impact parameter, b , which is the transverse distance between the centers of the colliding heavy-ions. It is defined such that $b = 0$ for central collisions, see Fig. 35.




FIG. 35: Left panel shows a side view of two relativistic heavy ions colliding. Right panel shows a view along the beam axis, where the cross-hatched almond-like overlap region is indicated. The reaction plane for a particular collision is the plane defined by the impact parameter, b , and the beam axis (z).

In most physics analyses of heavy ion collision data at highly relativistic energies, the impact parameter is not considered particularly useful in characterizing the important influence of geometry on the outcome of a

given interaction. Instead, two parameters which quantify the critical distinctions are used; namely the number of participating nucleons, N_{part} , and the number of binary nucleon-nucleon collisions, N_{coll} . In defining these variables, two important assumptions are made. First, since the collision duration at such high energies is very short compared to the typical time-scale for nuclear rearrangement or movement of nucleons within the nucleus, it is assumed that only the nucleons in the overlap region (the cross-hatched area in the right panel of Fig. 35) experience any substantial interactions (i.e. participate) in the collision. Second, the collisions suffered by a given nucleon as it traverses the other nucleus may not be distinct sequential events, and thus it may be most meaningful to simply count the total number of collisions. For observables such as elliptic flow which are sensitive to the shape of the initial overlap region, a third parameter, namely the spatial asymmetry derived from the impact parameter and the radii of the colliding nuclei, can be used.




FIG. 36: The left panel shows N_{part} and N_{coll} vs. impact parameter, b , for Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV. The right panel shows ν , the average number of collisions N_{coll} divided by the average number of participant pairs versus N_{part} for Au+Au at a variety of beam energies.

In determining the number of participating nucleons, or equivalently the number of nucleons which interact, only those which are struck by nucleons from the other nucleus (as opposed to ones which were hit only in secondary scatterings) are counted. This is the same quantity as “wounded nucleons” introduced by Białas, Bleszyński and Czyż [143]. N_{part} depends on the collision geometry and is typically calculated using a Glauber model of the collision. The key ingredients in this calculation are (1) nucleons are distributed according to a nucleon density function (e.g. Woods-Saxon), (2) nucleons in each nucleus travel in straight lines through the colliding system and (3) nucleons interact according to the inelastic cross section, σ_{NN} , as measured in proton-proton collisions. For the energies at RHIC, the values assumed were $\sigma_{NN}=33, 36, 41$, and 42 mb for $\sqrt{s_{NN}}=19.6, 62.4, 130$, and 200 GeV, respectively. See the open circles in the left panel of Fig. 36 for an example of the results of such a model calculation relating N_{part} and impact parameter for Au+Au collisions at one of the RHIC energies. The number of participants is usually assumed to

have a strong influence on the bulk properties of particle production but it is shown in the physics sections of this paper that N_{part} (or $N_{part}/2$) provides a convenient benchmark to study the effects of the collision geometry on many measured experimental quantities.

As introduced above, N_{coll} denotes the number of binary nucleon-nucleon collisions in a heavy ion reaction. As in the calculation of N_{part} , only primary collisions, i.e. those occurring along the straight-line trajectory of nucleons through the opposing nucleus, are counted. This quantity can also be calculated in a Glauber model, with typical results being shown as closed circles in the left panel of Fig. 36. Hard scattering (i.e. large momentum transfer) processes are expected to scale as N_{coll} . For symmetric A+A collisions, simple geometrical arguments imply that N_{coll} would scale as roughly $A^{4/3}$. Thus, for collisions of more than two participants, the number of binary nucleon-nucleon collisions is larger than the number of participants, with the difference increasing dramatically for smaller impact parameters.

One possibly important aspect of centrality in heavy ion collisions which goes beyond the simple increase in the number of participants or collisions is shown in the right panel of Fig. 36. There, the number of collisions is divided by the number of participating pairs to derive the average number of collisions suffered by each participant. A similar parameter, typically denoted ν and calculated from $\nu = (A\sigma_{pp})/\sigma_{pA}$ where the σ 's are inelastic cross-sections, is commonly used to characterize centrality or target dependences of observables in p+A collisions [91]. In nucleus-nucleus collisions, the calculated average number of collisions per participant varies by a large factor as a function of centrality and also has some dependence on energy due to the varying nucleon-nucleon cross section.

2. Particle Characterization

In describing the trajectories of particles emitted in heavy ion collisions, a distinction is typically made between longitudinal (i.e. along the beam) and transverse motion. The former may reflect some remnant of the original motion of the beam while the latter is largely generated in the interaction. The physics variable typically associated with the longitudinal motion is rapidity, denoted y and defined as $y = \frac{1}{2} \ln((E + p_{\parallel})/(E - p_{\parallel})) = \ln((E + p_{\parallel})/m_{\tau})$ with E and p_{\parallel} being the total energy and the component of the particle's momentum along the beam, respectively, and m_{τ} being the transverse mass. Rapidity has the important property of being additive in Lorentz transformations from one reference frame to another which differ by velocity along the beam. Thus, the shape of the distribution of any quantity plotted versus rapidity is the same in any frame. Unfortunately, it is frequently difficult to experimentally determine the particle identification, or in some cases even the momentum itself, necessary to calculate rapidity. In such instances, it is common to replace rapidity with pseudorapidity, de-

noted η and defined as $\eta = -\ln(\tan(\theta/2))$, where θ is the polar angle to the beam axis. For particles whose total momentum is large compared with their mass, i.e. for particles whose velocity is close to the speed of light ($\beta = v/c \approx 1$), the two measures are close to identical. Since the produced particles are typically dominated by pions whose transverse momentum alone averages a few hundred MeV/c or more, the use of pseudorapidity is a quite reasonable approximation. A variable frequently used in elementary collisions is the Feynman x_F variable given by the ratio of the momentum along the beam to the maximum possible value, $x_F = p_{\parallel}/p_{\parallel}^{max}$.

Another aspect of the distributions as a function of longitudinal velocity that proves to be very interesting is the comparison of distributions at a variety of beam energies but viewed in the rest frame of one of the projectile particles. For distributions as a function of rapidity, this can be done exactly and trivially by simply subtracting the rapidity of the beam from the rapidity of each particle. In the case of pseudorapidity distributions, the transformation is not exact but a reasonably close approximation is found using the shifted pseudorapidity, denoted η' and defined as $\eta' = \eta - y_{beam}$, where η is the pseudorapidity of a particle and y_{beam} is the beam rapidity. The quantity y_{beam} , which is given by $\frac{1}{2} \ln((E+p)/(E-p)) = \ln((E+p)/M)$ with E , p , and M being the energy, momentum, and mass of the beam, respectively, is equal to 3.04, 4.20, 4.94, and 5.37 for $\sqrt{s_{NN}}=19.6, 62.4, 130$, and 200 GeV, respectively. Fermi motion of 300 MeV/c would spread them out in rapidity by typically ≈ 0.3 units in rapidity.

The transverse motion is most often characterized using simply the component of the momentum, denoted p_T , that is perpendicular to the beam axis. Occasionally, the so-called transverse mass, $m_T = \sqrt{p_T^2 + m_0^2}$, is used where m_0 is the rest mass of the particle. The use of this more complicated variable is motivated by its appearance as the natural scaling parameter for particles emitted by a thermal source. It can also be used to combine energy, transverse momentum, and rapidity of a particle via the identity $E = m_T \cosh(y)$.

The various particle characterization variables can be related using the following identities:

$$p_{\parallel} = m_T \sinh(y) = p_T \sinh(\eta)$$

For relativistic beam energies, $p_{\parallel} \approx \frac{\sqrt{s}}{2} x_F$ and for y larger than about 1–2, $\sinh(y) \approx \frac{e^y}{2}$ so that:

$$\eta' \equiv \eta - y_{beam} \approx \ln(x_F) - \ln\left(\frac{p_T}{M}\right)$$

$$y' \equiv y - y_{beam} \approx \ln(x_F) - \ln\left(\frac{m_T}{M}\right),$$

where M is the nucleon mass.

In the case of jets emitted in e^+e^- annihilation, the motions of individual particles along and transverse to

the beam are not the interesting quantities. Instead, distributions are characterized by the trajectories of particles relative to the jet direction, the so-called thrust axis. Since data exist most frequently in the form of unidentified charged particles, the motion along the thrust axis is traditionally defined using y_T , the rapidity calculated using the momentum parallel to the jet direction and assuming the pion mass. The required shift to compare different beam energies in a common frame, as was done for y' or η' , is not intuitively obvious. In this paper, the somewhat arbitrary choice was made to replace y_{beam} in the formulas above with y_{jet} which is the rapidity calculated using the center-of-mass energy combined with the assumption of the proton mass. Therefore, the same shift was used in both e^+e^- and $p+p$ at the same \sqrt{s} .

3. Notation for Observables

The most basic observable characterizing particle production is the total number of particles emitted. Two experimental hurdles complicate the extraction of this number from the data. The first is that only charged particles are easily detected. Although assumptions can be made concerning the ratio of charged and neutral particles, the multiplicity data is almost always presented in terms of the number of charged particles. Adjustments for the number of unobserved neutrals is typically only done when needed in a specific calculation, for example in the discussion of the energy density presented in Sect. II A. The notation $N_{ch}(A+B)$ is used to denote the total charged particle yield, integrated over all phase space, in collisions of species A with species B. To date, PHOBOS has measured $N_{ch}(d+Au)$ at a variety of centralities for one center-of-mass energy and $N_{ch}(Au+Au)$ over a broad range of both centrality and beam energy. Note that in all cases, the multiplicity is defined to be “primary”, i.e. those particles emitted in the initial interaction. Corrections are applied to the data to remove all other “secondary” particles, which are created in weak or electromagnetic decays of primary particles and interactions of primary particles with material in the detector. The second complication in extracting total numbers is that no detector can be fully hermetic, i.e. capable of detecting every single particle emitted. As a result, it is always necessary to measure distributions of particles and extrapolate into the unmeasured regions.

Because the PHOBOS multiplicity detector measures only the emission angle of charged particles, the extracted distribution is the number of charged particles per unit pseudorapidity, denoted $dN_{ch}/d\eta$. The experimental layout is designed to minimize the amount of material between the collision vertex and the active elements and therefore the cut-off in transverse momentum is low and the losses of particles with low p_T are small. The correction for secondary particles which are added to the total by decays or interactions in the material are typically much larger than the correction for particles

that are lost. In addition, the very broad coverage in η provided by the PHOBOS setup results in a relatively small extrapolation for particles emitted at small angles with respect to the beam. Thus, PHOBOS can provide information about $dN_{ch}/d\eta$ and N_{ch} which are unique at RHIC. As mentioned above, it is also interesting to study particle distributions shifted into the rest frame of one of the projectiles. The shifted distribution, $dN_{ch}/d\eta'$, can be used as a measure of the charged particle pseudorapidity density as viewed in the rest frame of one of the colliding nuclei although one should keep in mind that such a shift in principle is associated with a small distortion of the distributions.

As is discussed in the main body of this paper, the particle density is highest near y or η of zero and therefore it is generally assumed that the potential for creation of any new state of matter is also highest in that region. As a result, the properties of observables “near mid-rapidity” are of particular interest. For the mid-rapidity multiplicity distribution, the range chosen is ± 1 unit in η so the pseudorapidity distribution is averaged over this range to generate $dN_{ch}/d\eta|_{|\eta|\leq 1}$.

In cases where the momentum and angle of the particles are measured, distributions in both transverse momentum and rapidity (or pseudorapidity in cases without particle identification) can be generated. The transverse distributions are commonly presented in a form which is Lorentz invariant given by $Ed^3\sigma/dp^3$, with E and p being the total energy and vector momentum of the particle, respectively. Since the interesting quantity is typically the number of particles in a given event, i.e. the distribution that integrates to give multiplicity, this is more commonly expressed as invariant yield Ed^3N/dp^3 . When integrating over all orientations of the reaction plane, azimuthal symmetry can be assumed and the differential momentum volume can be expressed in cylindrical coordinates as $dp^3 = 2\pi p_T dp_T dp_{||}$. Furthermore, the component of the momentum parallel to the beam can be transformed using $dp_{||} = Edy$ where y is the rapidity, resulting in the final form $d^2N/2\pi p_T dp_T dy$. When using transverse mass, the transformation is trivial since $p_T dp_T = m_T dm_T$ and only the horizontal axis changes in the distributions. In cases without particle identification, rapidity is approximated by pseudorapidity, yielding $d^2N/2\pi p_T dp_T d\eta$.

When comparing transverse momentum distributions for more complicated systems to data from proton-proton collisions, one could simply take the ratio of the two distributions as a function of p_T to study the change in magnitude or shape. This ratio is called the nuclear modification factor since it is a measure of the modification of the properties of the emitted particles resulting from the presence of the nucleus in the interaction. In order to test specific theories of how the yield should scale, the standard procedure is to normalize the A+A (or, equivalently, scale the p+p) data by some factor. The resulting ratio comparing collisions of species A with species B to

p+p is typically denoted R_{AB} defined as

$$R_{AB} = \frac{1}{Norm} \frac{dN_{A+B}/dp_T}{dN_{p+p}/dp_T} = \frac{1}{N_{coll}} \frac{dN_{A+B}/dp_T}{dN_{p+p}/dp_T}$$

The most common normalization, and the one usually indicated by the simple notations R_{AA} , R_{dAu} , etc. is N_{coll} as shown in the right formula above. This arises from the interest in studying the behavior of high transverse momentum particles and the belief that the such “hard” processes should scale with the number of binary nucleon-nucleon collisions. Analysis by the PHOBOS collaboration has demonstrated that the number of pairs of participating nucleons is often the more appropriate scaling variable. To avoid confusion, ratios using this latter normalization are denoted

$$R_{AB}^{N_{part}} = \frac{1}{N_{part}/2} \frac{dN_{A+B}/dp_T}{dN_{p+p}/dp_T}$$

Note that a p+p collision has one pair of participants. This normalization will be generically referred to as the number of participant pairs even in asymmetric collisions.

It is frequently of interest to study the evolution of the shape and magnitude of these distributions as a function of centrality for nucleus-nucleus collisions. The most direct display of this evolution involves dividing data from one centrality bin by that from a different bin. In this case, both distributions need to be suitably normalized. The notation PC (CP) is used for ratios where peripheral (central) data is divided by central (peripheral). The PHOBOS collaboration has recently advocated the use of R_{PC} since different experiments have different reach in centrality and the central data typically have significantly smaller statistical and systematic errors. In keeping with the convention described above, the definitions with the different normalizations are

$$R_{PC} = \frac{N_{coll}^{central} dN_{A+B}^{periph}/dp_T}{N_{coll}^{periph} dN_{A+B}^{central}/dp_T}$$

and

$$R_{PC}^{N_{part}} = \frac{N_{part}^{central} dN_{A+B}^{periph}/dp_T}{N_{part}^{periph} dN_{A+B}^{central}/dp_T}$$

Using an event-by-event measurement of the orientation of the two colliding nuclei, the study of particle distributions can be extended to include a third coordinate, namely the azimuthal angle. In relativistic heavy ion collisions, the generic terms “directed flow” and “elliptic flow” are used for the measurement of anisotropy in the azimuthal distributions of particles relative to the reaction plane. The reaction plane for a particular collision is the plane defined by the impact parameter, b , and the beam axis (b and z in Fig. 35). In flow analyses, the distribution of particles in the azimuthal angle, ϕ , (always taken relative to the reaction plane for a particular collision) is measured and expressed in terms of a Fourier expansion, $dN_{ch}/d\phi = N_0(1 + 2v_1\cos(\phi) + 2v_2\cos(2\phi) + \dots)$.

The amplitude of the first ϕ -dependent term, v_1 , is called directed flow. Elliptic flow is the name given to the amplitude of the second term of the Fourier expansion, v_2 . This latter anisotropy in the form of a variation in particle yield in momentum space is believed to result from the non-spherical shape in position space of the initial collision volume (see the cross-hatched region in the right panel of Fig. 35).

Moving beyond single particle distributions, additional information can be obtained by studying the correlations of particles. In heavy ion collisions, the most common multi-particle observable studied is the HBT correlation, named for Hanbury-Brown and Twiss who pioneered an analogous technique for studying the size of objects in astronomy [123, 124]. The procedure is most often applied to pairs of like-sign pions and depends on the quantum mechanical connection between separation in coordinate and momentum space for identical particles. The data are presented as the ratio of the distribution of pairs in some relative-momentum variable divided by a distribution which matches the correct occupancy of the two-particle phase space but which does not contain the effects of the two-particle correlation. This normalization is obtained by pairing particles found in different events which have been matched for centrality and other event-characterization variables. The resulting correlation functions can be fit using a variety of parameterizations of the source distributions. From such parameterizations, information about the spatio-temporal extent of the emission source can be extracted. One commonly used system is the so-called Bertsch-Pratt coordinates [144–146]. For a given pair of identical particles with average momentum k , the coordinates are: longitudinal (R_l) along the beam direction (z), outwards (R_o) in the (z, k) plane perpendicular to z , and sideways (R_s) perpendicular to the other two directions. The Yano-Koonin-Podgoretskii parameterization also includes spatial parameters for the longitudinal and transverse sizes of the source as well as parameters describing the duration and longitudinal velocity of the source [147, 148].

APPENDIX C: TECHNIQUES FOR DETERMINING CENTRALITY

As briefly discussed in Appendix B, determining the centrality of a heavy-ion collision is extremely important for event characterization. Knowing the centrality provides a geometrical scale for use in any studies of the underlying collision dynamics and affords the possibility of a more meaningful comparison to “baseline” data from elementary proton or electron collisions. The primary event centrality in PHOBOS is determined by utilization of signals from the Paddle scintillator counters as well as the Octagon and Ring silicon detectors, all of which are sensitive to charged particle multiplicities in various regions of phase-space. These signals, through bins in the percentage of total cross section, provide a

measure of centrality. The validity of this technique is based on the experimental observation of a strong correlation between the charged particle multiplicity signals in, for example, the Paddle scintillator counters and neutral beam “remnants” (neutrons) as measured in the Zero-Degree-Calorimeters (ZDCs). This correlation is shown in Fig. 37 for $\sqrt{s_{NN}}=200$ GeV Au+Au collisions at PHOBOS.




FIG. 37: Correlation between spectator neutrons measured in the PHOBOS ZDC's (ZDC Sum) and charged particle multiplicity measured in the Paddle counters (Paddle Mean) for $\sqrt{s_{NN}}=200$ GeV Au+Au collisions.

The specific methods developed within PHOBOS to determine centrality depend on both the collision species (Au+Au versus d+Au) and the collision energy. The technique is to associate an experimentally measured signal to a well-defined centrality related variable, such as the number of participating nucleons, N_{part} . For this technique to be meaningful, a monotonic relation must exist between the multiplicity signals in the chosen region of phase-space and N_{part} . This assumption is justified by the experimental correlation shown in Fig. 37 (the remnant neutrons are anti-correlated with N_{part} for the 50% most central collisions). Additional evidence for the validity of this technique has been obtained using extensive Monte Carlo (MC) studies using event generators (such as HIJING, AMPT, RQMD, and Venus) and a full GEANT simulation of the PHOBOS detector. An outline of some of these techniques follows.

1. Centrality determination in Au+Au collisions

There are four main considerations that must be addressed in the course of determining the event centrality; the event selection, detection efficiency, choice of phase-

space region to utilize, and the event generator simulations to extract N_{part} .

The initial event selection must cleanly identify and separate true Au+Au collisions from numerous background sources, such as beam-gas interactions, while simultaneously providing the smallest possible bias on the resulting data set. In PHOBOS this was accomplished by using a combination of energy and time signals from the Paddle counters and the ZDCs. A selection of events with less than 4 ns time-difference between the two Paddle signals was combined with cuts on the ZDC individual and summed timing signals. Additional logic ensured no loss of very central events that have a high Paddle signal and corresponding few numbers of spectator neutrons available to hit the ZDCs. This selection provided a basic “valid collision” definition for Au+Au collisions at $\sqrt{s_{NN}}=62.4, 130$ and 200 GeV. For the lowest energy Au+Au collision of $\sqrt{s_{NN}}=19.6$ GeV, the ZDC timing requirement had to be modified for a diminishingly low efficiency for measuring the lower energy neutrons.




FIG. 38: Illustration of the detection efficiency determination in Au+Au collisions using a comparison between Monte Carlo (MC) simulation and data for the number of Paddle slats hit. Data are shown for 200 GeV Au+Au collisions, with the same technique used for 62.4 and 130 GeV Au+Au collisions.

The detector efficiency was determined for two “minimum-bias” trigger configurations of at least one (two) hits in each scintillator Paddle counter array. For both configurations, the trigger system incurred a loss of peripheral events that had to be accounted for before percentage of total cross section bins could be correctly fashioned. The fraction of lost peripheral events was determined using comparisons of the total number of Paddle slats hit in both data and the full MC simulations (see Fig. 38). This analysis yielded a detection efficiency of 97% and 88% for the two trigger configurations, respectively, for Au+Au collisions at $\sqrt{s_{NN}}=200$ GeV.

Using the valid collision event selection criteria outlined above and the deduced trigger detection efficiency, the next task is to find an appropriate experimental quantity for use in determining the event centrality. For Au+Au collisions, a consistent centrality determination was found to be relatively independent of the choice of phase-space region used, as long as the chosen region in pseudorapidity (η) was within the range of substantial particle multiplicity. The signal from the Paddle counters, with a phase-space coverage of $3.2 < |\eta| < 4.5$ (region (b) of Fig. 39), worked well as a centrality measure for collision energies of $\sqrt{s_{NN}}=62.4, 130$ and 200 GeV.




FIG. 39: Pseudorapidity density distributions from $\sqrt{s_{NN}}=200$ (light, top band) and 19.6 (dark, bottom band) GeV Au+Au collisions, for the most central 25% of the cross section [24]. The boxed areas (a–d) illustrate the separate phase-space regions in pseudorapidity used in the centrality determination for each collision energy. Region (b) illustrates the pseudorapidity coverage of the Paddle scintillator counters, and the other regions were developed for centrality determination using the Octagon silicon detector.

For the lowest energy of 19.6 GeV, new phase-space regions had to be chosen due to a reduction in the monotonicity between the multiplicity signals in the Paddle counters and both the number of spectator neutrons seen in the ZDCs and N_{part} , as determined from MC simulations. In addition, the Paddles are traversed by significantly fewer particles at 19.6 than 200 GeV (see dark grey band in Fig. 39 or the bottom panel of Fig. 1) and, consequently, a different phase space region had to be defined. In order to create a centrality measure at 19.6 GeV similar to that obtained from the Paddles at 200 GeV, the Paddle pseudorapidity range was scaled down to a smaller region by the ratio of beam rapidities $y_{19.6}^{beam}/y_{200}^{beam} = 0.563$ (region (d) of Fig. 39). The resulting η region, $1.8 \leq |\eta| \leq 2.5$, lies wholly within the Octagon silicon detector coverage of $|\eta| \leq 3.2$ for colli-

sions which occur within ± 10 cm of the nominal vertex position. Thus, the charged particle multiplicity measured in the region (d) was used as a centrality measure for 19.6 GeV and allowed for a direct centrality comparison to the original Paddle-based method at 200 GeV. Additional centrality measures were developed at both energies, in phase space regions close to mid-rapidity, which used the multiplicity signals of charged particles traversing the Octagon in the pseudorapidity regions (a) and (c), where region (c) is scaled by a factor of 0.563 compared to region (a). This technique of matching centrality regions allowed direct comparisons of mid-rapidity and away from mid-rapidity centrality determinations across a factor of ten in collision energy. Also, both pseudorapidity regions have been found to have very different rates of particle production and intrinsic dependences on N_{part} . By utilizing these two independent regions, the assumption that the centrality measure and N_{part} have to be only monotonic and not necessarily linear can be explicitly tested. An insignificant difference was found when analyzing data with both methods, at both energies [26].




FIG. 40: Charged particle multiplicity signal distributions used in the four centrality determination pseudorapidity regions (a–d) taken from Fig. 39. Black histograms are data and the grey distributions are MC simulations. All data are shown for a restricted collision vertex of $|z| \leq 10$ cm, and thus have an additional inefficiency for low multiplicities as evident from the figures where the data falls below the (unbiased) MC simulations for peripheral events.

Use of the Octagon silicon detector signals as a centrality measure introduces an additional complication not present for the Paddle counters. The precise vertex position of each event is required for the merging and angle correction of valid hits in the Octagon. PHOBOS has developed several techniques to determine the primary collision vertex, including use of the Vertex detector and

straight-line tracks in the first six planes of the Spectrometer. However, due to the requirement of any of these valid vertices, the resulting data set is not only biased by the intrinsic trigger efficiency, but also by the vertex reconstruction efficiency. Additional inefficiencies are introduced for low multiplicity events and this fact is the primary reason that PHOBOS has, thus far, only published events for the top 50% of cross section for Au+Au data, where there are no inefficiencies present. Despite these additional complications, exploiting the Octagon detector signals as a centrality measure greatly expands the available phase space for centrality determination. As shown in Fig. 40, a good match between the data and MC simulations in all regions of phase-space shown in Fig. 39 gives confidence in the validity of the procedure.

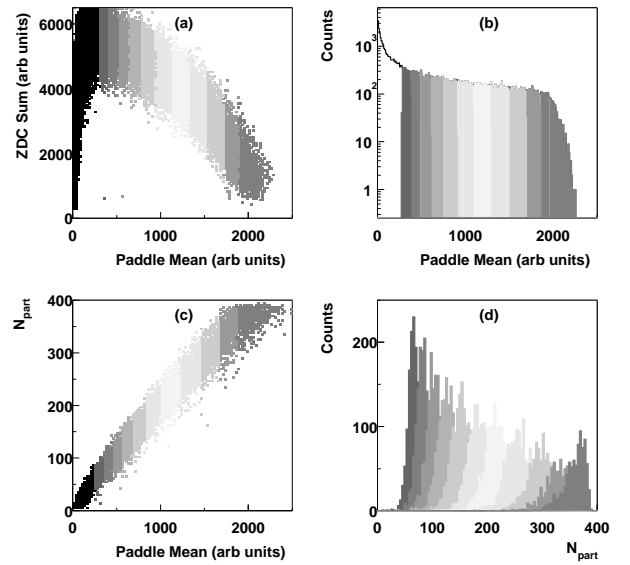


FIG. 41: Illustration of how the centrality of a heavy-ion Au+Au collision is defined. Only the top 50% of cross section, where there is 100% detection and vertex reconstruction efficiency, is used. Panel (a) shows the experimental correlation between the charged particle multiplicity signals in the Paddle counters (Paddle Mean) and the signals in the ZDC's from spectator neutrons. The shaded bands represent bins in percentile of cross section cut on the Paddle Mean signal. Panel (b) is a projection of (a) onto the Paddle Mean axis. Panel (c) shows a corresponding MC calculation where a monotonic relation is observed between the Paddle Mean signal and N_{part} , the number of participating nucleons. From this correlation, the average N_{part} (see panel (d)) can be extracted for each bin in percentile of cross section.

Once the choice of phase-space region for the centrality determination is made and the corresponding efficiency is determined, the resulting multiplicity related distribution can be divided into percentile of total cross section bins, as illustrated in Fig. 41, panels (a) and (b). Comprehensive MC simulations of these signals, that include Glauber model calculations of the collision geometry, allow the estimation of N_{part} for a cross section bin, as il-

illustrated in Fig. 41, panels (c) and (d). The most central collisions ($b \sim 0$, see Fig. 35) will have the largest number of participants with the obvious upper limit of $N_{part}=394$ for a “perfectly central” Au+Au collision, where all nucleons interact.

Systematic uncertainties on the extracted values of N_{part} were determined with extensive MC simulations that included possible errors in the overall detection efficiency and also utilized different types of event generators. The uncertainty on the deduced N_{part} increased from $\sim 3\%$ for central collisions to $\sim 9\%$ for mid-peripheral.

2. Centrality determination in d+Au collisions

The centrality determination for d+Au collisions at $\sqrt{s_{NN}}=200$ GeV incorporates the same considerations necessary for Au+Au collisions, however the details of the analysis techniques differ.

The initial definition of a ‘valid collision’ took on a different form, as the most basic event selection could not be defined cleanly from the timing signals of the Paddles. The asymmetric nature of the collision system resulted in very different particle multiplicities impinging on the two symmetrically located Paddle trigger counters, which caused an overall reduction in the timing resolution compared to that of the Au+Au collision system. Previous requirements of a ZDC detector timing coincidence were also no longer possible due to the unacceptable bias that would be imposed on the dataset. To ensure that the events analyzed were real collisions occurring in a usable proximity to the detector, a reconstructed collision vertex was required.

Lower total multiplicities precluded the high resolution track-based vertex reconstruction algorithms developed for Au+Au collisions ($\sigma_z \leq 0.04$ cm for central collisions) as it became increasingly inefficient. A more efficient, but less accurate ($\sigma_z \sim 0.8$ cm for central collisions) vertex reconstruction method based on global averaging techniques across the entire Octagon was created. This selection coupled with the intrinsic triggering of the system was estimated to have an overall efficiency of 83% for d+Au collisions at 200 GeV. This high efficiency data set was used for more global physics analyses, such as the centrality dependence of the $dN_{ch}/d\eta$ distribution.

In addition to the Octagon based vertex determination, a new on-line fast vertex position derived from the T0 Cerenkov counters (see Appendix A) was developed for the d+Au collision data. This T0 time-difference based method was utilized both as a primary trigger for some data sets as well as, when fully calibrated, an additional off-line vertex finding algorithm to aid the new Octagon-based algorithm discussed above in order to provide a very clean event selection. An additional benefit of on-line vertex triggering from the T0 detectors was the enhancement of the fraction of data occurring at the center of the detector ($|z| < 20$ cm). A Paddle-triggered dataset

in d+Au allowed data to be written for collisions occurring within approximately 2 m. The additional (T0) requirement forced a larger bias on the data than the Paddles and Octagon vertices alone and further reduced the overall trigger+vertex efficiency to 49%, but resulted in a much higher fraction of usable data that provided significantly improved statistics necessary for many Spectrometer based analyses.




FIG. 42: Reconstructed MC simulated pseudorapidity distributions for d+Au collisions at 200 GeV for peripheral (panel a) and central (panel b) collisions where the centrality definition is taken from different regions of phase-space (see text for discussion). MC simulations shown utilized the HIJING event generator coupled to a complete GEANT simulation of the PHOBOS detector.

The more significant challenge in the d+Au data analysis was to extract the centrality dependence of various physics analyses without the centrality measure itself directly influencing the outcome as a result of strong auto-correlations. This issue is not a key factor when measuring quantities for a minimum-bias configuration [149], but it becomes a major consideration for any detailed studies requiring a centrality definition. In these cases, unlike Au+Au collisions, the centrality determination for d+Au collisions was found to be strongly dependent on the choice of phase-space region utilized in the analysis. This fact is illustrated with MC simulated data in Fig. 42, where strong auto-correlation biases are seen in the reconstructed pseudorapidity distributions for four of the five different centrality methods explored. Specifically a suppression of mid-rapidity yields ($|\eta| < 3$) in the reconstructed spectrum for peripheral collisions is observed Fig. 42a. The exact opposite effect is observed for central collisions, i.e. enhancement at mid-rapidity (Fig. 42b). This study utilized five different centrality definitions that each covered different regions of pseudorapidity; the Octagon detector ($E_{Oct}, |\eta| \leq 3.0$), the Ring detectors ($E_{Ring}, 3.0 \leq |\eta| \leq 5.4$), the combined coverage of Octagon and Rings ($E_{Tot}, |\eta| \leq 5.4$), the deuteron direction ($E_{dHem}, 0.5 \leq \eta \leq 5.4$) and the gold direction ($E_{AuHem}, -5.4 \leq \eta \leq -0.5$). Both HIJING and AMPT based MC simulations indicated that a centrality measure based on the signals in the Ring counters provided the least bias on the measurement.

Additional support for using the MC based simulations to select the best centrality measure is given in Fig. 43. Ratios of the reconstructed $dN_{ch}/d\eta$ distributions obtained from four centrality measures relative to




FIG. 43: (NOTE: Rachid will make a new figure of this) Ratios of reconstructed $dN_{ch}/d\eta$ distributions for both Data and MC simulations using different centrality measures, each of which is selecting on the same percentile of central collisions. The good agreement in these ratios gives confidence that the MC simulations are providing a good basis on which to study the effects of biases created in the data that result from using different regions of pseudorapidity for the centrality determination.

that obtained using E_{Ring} variable are compared between data and MC simulations. These ratios are found to be in very good agreement. This information, which is based on data and MC simulation independently, provides the necessary confidence that using the Ring detectors for the centrality measure will provide the most accurate experimental result. It is important to point out that this study only provided guidance as to the choice of the Rings for the experimental centrality measure, and the final experimentally measured $dN_{ch}/d\eta$ distributions do not rely on the details of the MC simulation.

The choice of the Ring detectors for use in the centrality determination for d+Au collisions, along with the extracted efficiency, allows the creation of centrality bins based on percentage of cross section. For d+Au, a centrality determination is desired over the entire range of peripheral to central collisions. Thus corrections must be made to both the location of the E_{Ring} bins and the extracted N_{part} values to properly account for the inefficiencies in detecting peripheral collisions. These corrections were made based on extensive MC simulations using both the HIJING and AMPT event generators. For the case of the T0 triggered dataset, the centrality determination was reanalyzed in terms of the efficiency as-

sociated with each cross section bin and the associated N_{part} . The additional requirement of hits (particles) in the T0 region of phase space serves to push the average N_{part} higher, with the largest shifts for lower centrality classes. An example of the resulting centrality bins on




FIG. 44: Illustration of how the centrality is defined for a d+Au collision. The entire cross section range is used in the analysis. The shaded bands represent bins in percentile of cross section cuts on the multiplicity signals in the Ring detectors. The data is shown for an online vertex (T0) restricted data set (see text).

the Ring signal distributions for four-bins of cross section are shown in Fig. 44. The decreasing efficiency for more peripheral collisions is immediately evident.

Systematic uncertainties on the deduced average N_{part} values for each percentile bin of cross section were determined with extensive additional simulations. In these studies, the N_{part} distribution taken directly from Glauber model calculations was matched to the measured centrality related variable, i.e. E_{Ring} , distribution from data, and the average corresponding N_{part} was extracted for each centrality bin. Many different effects, including various types of detector resolution smearing, possible non-linear dependencies of the measured centrality variable on N_{part} and different deuteron wave functions, were included and the analysis repeated. These studies showed that the mean value of N_{part} estimated from the full HIJING (or AMPT) + GEANT detector simulation was reasonable and the systematic error on N_{part} reaches $\sim 30\%$ for the most peripheral centrality bin, where the overall bias is greatest.

[1] H. Kastrup, P. Zerwas, editors, *QCD 20 years Later*, World Scientific, Singapore (1993).

[2] S. Eidelman *et al.*, Phys. Lett. B**592**, 1 (2004).

- [3] D. J. Gross and F. Wilczek, Phys. Rev. Lett. **30**, 1343 (1973).
- [4] H. Politzer, Phys. Rev. Lett. **30**, 1346 (1973).
- [5] J. C. Collins and M. J. Perry, Phys. Rev. Lett. **34**, 1353 (1975).
- [6] E. V. Shuryak, Phys. Reports **61**, 71, (1980).
- [7] “Statistical Mechanics of Quarks and Hadrons”, H. Satz, ed., (North-Holland, 1981).
- [8] D. J. Gross, R. D. Pisarski, and L. G. Yaffe, Rev. Mod. Phys. **53**, 43 (1981).
- [9] W. Marciano and H. Pagels, Phys. Reports **88**, 137 (1978).
- [10] R.D. Pisarski and F. Wilczek, Phys. Rev. D**29**, 338 (1984).
- [11] F. Karsch, Nucl. Phys. **A698**, 199 (2002).
- [12] T. Schäfer, arXiv:hep-ph/0304281 (2003).
- [13] M. G. Alford, Ann. Rev. Nucl. Part. Sci. **51**, 131 (2001).
- [14] K. Rajogopal and F. Wilczek, arXiv:hep-ph/0011333 (2000).
- [15] See the proceedings of the series of conferences “Quark Matter-International Conference on Ultra-Relativistic Nucleus-Nucleus Collisions”; J. Phys. G**30**, No. 8, S633 (2004); Nucl. Phys. **A715**, 1c (2003); *ibid.* **A698**, 1c (2002); *ibid.* **A661**, 1c (1999); *ibid.* **A638**, 1c (1998); *ibid.* **A610**, 1c (1996); *ibid.* **A590**, 1c (1995); *ibid.* **A566**, 1c (1994); *ibid.* **A544**, 1c (1992); *ibid.* **A525**, 1c (1991); *ibid.* **A498**, 1c (1989); Z. Phys. **C38**, 1 (1988); Nucl. Phys. **A461**, 1c (1987); “Quark Matter ’84”, ed. K. Kajantie (Springer Verlag, 1985); Nucl. Phys. **A418**, 1c (1984); *Quark Matter Formation and Heavy Ion Collisions*, Proceedings of the Bielefeld Workshop 1982, eds. M. Jacob and H. Satz (WSPC, Singapore, 1983); Phys. Reports **88**, 321 (1982); *Workshop on future relativistic heavy ion experiments*, R. Bock and R. Stock, eds., (Darmstadt, 1980).
- [16] Reference to the White Papers from the other three RHIC experiments
- [17] K. Kanaya, Nucl. Phys. **A715**, 233c (2003).
- [18] M. Plümer, S. Raha, and R. M. Weiner, Nucl. Phys. **A418**, 549c (1984).
- [19] E. V. Shuryak, arXiv:hep-ph/0405066.
- [20] F. Karsh *et al.*, Nucl. Phys. **A715**, 701 (2003).
- [21] C. T. H. Davies *et al.*, Phys. Rev. Lett. **92**, 022001 (2004).
- [22] C. R. Allton *et al.*, Phys. Rev. D**66**, 074507 (2002).
- [23] K. J. Eskola, Nucl. Phys. **A698**, 78c (2002).
- [24] B. B. Back *et al.*, Phys. Rev. Lett. **91**, 052303 (2003).
- [25] B. B. Back *et al.*, Phys. Rev. Lett. **88**, 022302 (2002).
- [26] B. B. Back *et al.*, Submitted to Phys. Rev. C Rapid Comm.; nucl-ex/0405027 (2004).
- [27] B. B. Back *et al.*, Phys. Rev. Lett. **85**, 3100 (2000).
- [28] B. B. Back *et al.*, Phys. Rev. Lett. **93**, 082301 (2004).
- [29] S. A. Bass *et al.*, Nucl. Phys. **A661**, 205 (1999).
- [30] K.J. Eskola, K. Kajantie, P. V. Ruuskanen and K. Tuominen, Nucl. Phys. **B570**, 379 (2000).
- [31] A. Krasnitz and R. Venugopalan, Phys. Rev. Lett. **86**, 1717 (2001).
- [32] D. Kharzeev and E. Levin, Phys. Lett. B **523**, 79 (2001).
- [33] D. Kharzeev, E. Levin, and L. McLerran, BNL preprint BNLNT-02/22; arXiv:hep-ph/0210332.
- [34] D. Kharzeev, E. Levin, and L. McLerran, Phys. Lett. **B561**, 93 (2003).
- [35] D. Kharzeev, Y.V. Kovchegov, and K. Tuchin, arXiv:hep-ph/0405045.
- [36] B. B. Back *et al.*, Phys. Rev. Lett. **91**, 072302 (2003).
- [37] B. B. Back *et al.*, Submitted to Phys. Rev. C (RC); arXiv:nucl-ex/0406017 (2004).
- [38] I. Arsene *et al.*, submitted to Phys. Rev. Lett.; arXiv:nucl-ex/0403005 (2004).
- [39] M.X. Liu *et al.*, arXiv:nucl-ex/0403047.
- [40] A. Accardi, arXiv:nucl-th/0405046.
- [41] K. Adcox *et al.*, Phys. Rev. Lett. **87**, 052301 (2001).
- [42] J. Adams *et al.*, arXiv:nucl-ex/0407003 (2004).
- [43] B. B. Back *et al.*, Phys. Lett. **B578**, 297 (2004).
- [44] B. B. Back *et al.*, Submitted to Phys. Rev. Lett.; nucl-ex/0401006 (2004).
- [45] S. S. Adler *et al.*, Phys. Rev. C**69**, 034909 (2004).
- [46] J. D. Bjorken, Phys. Rev. D **27**, 140 (1983).
- [47] B. B. Back *et al.*, Phys. Rev. Lett. **87**, 102301 (2001).
- [48] B. B. Back *et al.*, Phys. Rev. C**67**, 021901R (2003).
- [49] F. Becattini, J. Cleymans, A. Keränen, E. Suhonen, and K. Redlich, Phys. Rev. C**64**, 024901 (2001).
- [50] B. B. Back *et al.*, Accepted for publication in Phys. Rev. C Rapid Comm.; nucl-ex/0309013 (2003).
- [51] F. Karsch, J. Phys. G **30**, S887 (2004).
- [52] C. Bernard *et al.*, arXiv: hep-lat/0405029 (2004).
- [53] F. Becattini, M. Gazdzicki, A. Keränen, J. Manninen and R. Stock, Phys. Rev. C**69**, 024905 (2004).
- [54] B. B. Back *et al.*, Phys. Rev. Lett. **89**, 222301 (2002).
- [55] B. B. Back *et al.*, Submitted to Phys. Rev. C (RC); arXiv:nucl-ex/0407012 (2004).
- [56] K. H. Ackermann *et al.*, Phys. Rev. Lett. **86**, 402 (2001).
- [57] P. Huovinen, private communication of a calculation using the results of P. F. Kolb, P. Huovinen, U. W. Heinz, and H. Heiselberg, Phys. Lett. **500**, 232 (2001).
- [58] P. F. Kolb and U. Heinz, Hydrodynamical description of ultrarelativistic heavy-ion collisions. In Quark-Gluon Plasma III. (World Scientific, 2004); arXiv:nucl-th/0305084 (2003).
- [59] W. Busza, in *Particle Production in Highly Excited Matter*, edited by H. H. Gutbrod and J. Rafelski (Plenum Press, Proceedings of a NATO Advanced Study Institute on Particle Production in Highly Excited Matter, 1993) p. 149.
- [60] J. C. Collins, D. E. Soper, and G. Sterman, Nucl. Phys. **B261**, 104 (1985).
- [61] J. C. Collins, D. E. Soper, and G. Sterman, “Factorization of Hard Processes in QCD”, in “Perturbative QCD”, A.H. Mueller, ed., World Scientific Publ., (1989); Adv. Ser. Direct. High Energy Phys. **5**, 1 (1988).
- [62] K. Adcox *et al.*, Phys. Rev. Lett. **88**, 022301 (2002).
- [63] K. Adcox *et al.*, Phys. Lett. **B561**, 82 (2003).
- [64] B. B. Back *et al.*, Submitted to Phys. Rev. Lett.; nucl-ex/0405003 (2004).
- [65] I. Arsene *et al.*, Phys. Rev. Lett. **91**, 072305 (2003).
- [66] S. S. Adler *et al.*, Phys. Rev. Lett. **91**, 072303 (2003).
- [67] J. Adams *et al.*, Phys. Rev. Lett. **91**, 072304 (2003).
- [68] C. Adler *et al.*, Phys. Rev. Lett. **90**, 082302 (2003).
- [69] J. Adams *et al.*, arXiv:nucl-ex/0407007, submitted to Phys. Rev. Lett. (2004).
- [70] S. S. Adler *et al.*, Phys. Rev. Lett. **91**, 182301 (2003).
- [71] J. Adams *et al.*, Phys. Rev. Lett. **92**, 052302 (2004).
- [72] D. Molnar and S. A. Voloshin, Phys. Rev. Lett. **91**, 092301 (2003).
- [73] B. B. Back *et al.*, Submitted to Phys. Rev. Lett.; nucl-ex/0301017 (2003).
- [74] J. Klay, U.C. Davis PhD. Thesis (2001).
- [75] S. V. Afanasiev *et al.*, Phys. Rev. C**66**, 054902 (2002).

- [76] G. I. Veres, PhD Thesis, <http://na49info.cern.ch/cgi-bin/wwwd-util/NA49/NOTE?292> (2002).
- [77] A. E. Brenner *et al.*, Phys Rev D **26**, 1497 (1982).
- [78] M. Basile *et al.*, Phys. Lett. B **92**, 367 (1980).
- [79] M. Basile *et al.*, Phys. Lett. B **95**, 311 (1980).
- [80] B. B. Back *et al.*, Phys. Rev. Lett. **86**, 1970 (2001).
- [81] B. B. Back *et al.* dAu mult vs centrality. To be submitted.
- [82] I. G. Bearden *et al.*, arXiv:nucl-ex/0312023 (2003).
- [83] W. Busza and A. S. Goldhaber, Phys. Lett. B **139**, 235 (1984).
- [84] W. Busza and R. Ledoux, Ann. Rev. Nucl. Science **38**, 119 (1988).
- [85] G. J. Alner *et al.*, Phys. Rep. **154**, 247 (1987).
- [86] J. E. Elias *et al.*, Phys. Rev. D **22**, 13 (1980).
- [87] J. Whitmore, Phys. Rep. **10**, 273 (1974).
- [88] J. Whitmore, Phys. Rep. **27**, 187 (1976).
- [89] H. Boggild *et al.*, Ann. Rev. Nucl. Science **24**, 247 (1974).
- [90] W. Busza *et al.*, Phys. Rev. Lett. **34**, 836 (1975).
- [91] W. Busza, Acta. Phys. Pol. **B8**, 33 (1977).
- [92] B. B. Back *et al.*, Phys. Rev. Lett. **87**, 102303 (2001).
- [93] D. H. Brick *et al.*, Phys. Rev. D **41**, 765 (1990).
- [94] C. DeMarzo *et al.*, Phys. Rev. D **29**, 2476 (1984).
- [95] C. Halliwell *et al.*, Phys. Rev. Lett. **39**, 1499 (1977).
- [96] D. S. Barton *et al.*, Phys. Rev. D **27**, 2580 (1983).
- [97] H. Stenzel, ALEPH Collaboration, Contributed paper to ICHEP2000 (2000).
- [98] PYTHIA manual, T. Sjöstrand, Comp. Phys. Comm. **82**, 74 (1994). JETSET 7.4 is currently part of the PYTHIA code.
- [99] G. Arnison *et al.*, Phys. Lett. B **118**, 167 (1982).
- [100] C. Albajar *et al.*, Nucl. Phys. B **335**, 261 (1990).
- [101] G. Bocquet *et al.*, Phys. Lett. B **366**, 434 (1996).
- [102] J. Benecke, T. T. Chou, C. N. Yang, E. Yen, Phys. Rev. **188**, 2159 (1969).
- [103] T. T. Chou and C.-N. Yang, Phys. Rev. Lett. **25**, 1072 (1970).
- [104] R. P. Feynman, Phys. Rev. Lett. **23**, 1415 (1969).
- [105] R. P. Feynman, Photon-Hadron Interactions, (1972) pub. W.A. Benjamin Inc., p. 237-249.
- [106] T. Sjöstrand, P. Edén, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna, E. Norrbin, Comp. Phys. Comm. **135**, 238 (2001).
- [107] T. Sjöstrand, L. Lönnblad, S. Mrenna, P. Skands, arXiv:hep-ph/0308153.
- [108] W. Thomé *et al.*, Nucl. Phys. B **129**, 365 (1977).
- [109] G. J. Alner *et al.*, Z. Phys. C **33**, 1 (1986).
- [110] P. Abreu *et al.*, Phys. Lett. B **459**, 397 (1999).
- [111] B. Andersson, in *Proceedings VII International Colloquium on Multiparticle Reactions, Tutzing, 1976*, and references therein.
- [112] K. Gottfried, in *Proceedings V International Colloquium on Multiparticle Reactions, Uppsala, 1974*.
- [113] N. N. Nikolaev, Sov. J. Part. Nucl. **12**, 63 (1981) and references therein.
- [114] L. Stodolsky, in *Proceedings VI International Colloquium on Multiparticle Reactions, Oxford, 1975*.
- [115] I. Otterlund *et al.*, Nucl. Phys. B **142**, 445 (1978) and references therein.
- [116] S. Fredriksson *et al.*, Phys. Rep. **144**, 187 (1987) and references therein.
- [117] I. G. Bearden *et al.* arXiv:nucl-ex/0403050 (2004).
- [118] J. L. Klay *et al.* Phys. Rev. C **68**, 054905 (2003).
- [119] I. G. Bearden *et al.*, Phys. Lett. B **523**, 227 (2001).
- [120] I. G. Bearden *et al.*, Phys. Rev. Lett. **88**, 202301 (2002).
- [121] P. F. Kolb, Proceedings of the 17th Winter Workshop on Nuclear Dynamics, Park City Utah 2001. Acta Phys. Hung. New Ser. Heavy Ion Phys. **15**, 279 (2002); nucl-th/0104089.
- [122] B. B. Back *et al.*, Submitted to Phys. Rev. Lett.; arXiv:nucl-ex/0406021 (2004).
- [123] R. Hanbury-Brown and R. Q. Twiss, Phil. Mag. Ser. 7, Vol. 45, No. 366, 663 (1954).
- [124] R. Hanbury-Brown and R. Q. Twiss, Nature **178**, 1046 (1956).
- [125] D. Rischke and M. Gyulassy, Nucl. Phys. A **597**, 701 (1996).
- [126] T. Hirano and K. Tsuda, Phys. Rev. C **66**, 054905 (2002).
- [127] S. Soff, arXiv:hep-ph/0202240 (2002).
- [128] J. Adams *et al.*, Phys. Rev. Lett. **93** 012301 (2004).
- [129] S. S. Adler *et al.*, arXiv:nucl-ex/0401003 (2003).
- [130] B. B. Back *et al.*, arXiv:nucl-ex/0409001 (2004).
- [131] M. D. Baker, Proc. of the Eighteenth Lake Louise Winter Institute, Lake Louise, Alberta, Canada Feb. 2003; arXiv:nucl-ex/0309002 (2003).
- [132] T. Renk, arXiv:hep-ph/0404140; Submitted to Phys. Rev. C. (2004).
- [133] D. Teaney, arXiv:nucl-th/0301099 (2003).
- [134] H. Appelshäuser *et al.*, Eur. Phys. J. C **2**, 661 (1998).
- [135] B. B. Back *et al.*, Phys. Rev. C **65**, 031901R (2002).
- [136] B. B. Back *et al.*, Phys. Rev. C **65**, 061901R (2002).
- [137] D. Kharzeev and M. Nardi, Phys. Lett. B **507**, 121 (2001).
- [138] D. Kharzeev, E. Levin, and M. Nardi, arXiv:hep-ph/0111315 (2001).
- [139] B. B. Back *et al.*, Nucl. Inst. Meth. A **499**, 603 (2003).
- [140] B. B. Back *et al.*, Nucl. Instr. Meth. A **447**, 257 (2000).
- [141] B. B. Back *et al.*, Nucl. Phys. B **78** 245 (1999).
- [142] B. B. Back *et al.*, Nucl. Inst. Meth. A **419**, 549 (1998).
- [143] A. Białas, M. Bleszyński, and W. Czyż, Nucl. Phys. B **111**, 461 (1976).
- [144] S. Pratt, Phys. Rev. D **33**, 1314 (1986).
- [145] G. F. Bertsch, Nucl. Phys. A **498**, 173 (1989).
- [146] U. Heinz, Nucl. Phys. A **610**, 264 (1996).
- [147] F. Yano and S. Koonin, Phys. Lett. B **78**, 556 (1978).
- [148] M. Podgoretskii, Sov. J. Nucl. Phys. **37**, 272 (1983).
- [149] B. B. Back *et al.*, Accepted for publication in Phys. Rev. Lett.; nucl-ex/0311009 (2003).